

Quantitative X-ray Radiograph Analysis Jan/Feb 2006 Workshop Research Synopsis

Data-Driven Modeling and Analysis (DDMA) Team

Primary Contacts

Tom Asaki (CCS-2, asaki@lanl.gov)
Katharine Chartrand (CCN-8, kncx@lanl.gov)

Participants

Bill Allard (Duke University)
Erik Bollt (Clarkson University)
Patrick Campbell (T-7)
David Caraballo (Georgetown University)
Rick Chartrand (T-7)
David Dreisigmeyer (CCN-8)
Selim Esedoglu (University of Michigan)
Ousseini Lankoande (University of New Mexico, EES-11)
Chris Orum (D-1)
Robert Owczarek (EES-11)
Paul Rodriguez (T-7)
Pete Schultz (Clarkson University, T-7)
Matt Sottile (CCS-1)
Kevin Vixie (T-7)
Curt Vogel (Montana State University)

Other Contributors

Mark Abramson (Air Force Institute of Technology)
David Arathorn (Montana State University)
John Dennis Jr. (Rice University)
Kevin O'Reilly (Air Force Institute of Technology)
Brain Temple (X-1-SEC4)
Brendt Wohlberg (T-7)

Contents

1	Executive Summary	4
1.1	Motivation	4
1.2	Workshop Results	4
1.3	Discussion	6
2	Object Identification Within Radiographs	7
2.1	Transformation Discovery via Map-Seeking Circuit	8
2.1.1	Principles	8
2.1.2	Application	8
2.1.3	Example	8
3	Object Radiograph Extraction	12
3.1	Symmetry-Constrained Inpainting	13
3.1.1	Principles	13
3.1.2	Generalized Symmetry Discussion	15
3.1.3	Example	16
3.1.4	Implementation Details	17
4	Scattering Effects Removal	21
4.1	Total Variation Based Deconvolution	22
4.1.1	TV Overview	22
4.1.2	Method	23
4.1.3	Total Variation Descattering	23
4.1.4	Results	25
4.1.5	Conclusions	28
4.2	Heat-kernel Regularization and Inpainting by Linear Regression	33
4.2.1	Principles	33
4.2.2	Application	34
4.2.3	Examples	36
4.3	1D Smooth Function Inpainting	40
4.3.1	Principles	40
4.3.2	Application	41
4.3.3	Examples	41
4.4	n-Gaussians Model of Scattering Signatures	43
4.4.1	Principles	43
4.4.2	Discussion	44
4.4.3	Implementation	44
4.5	Biharmonic Inpainting	45
4.5.1	Concept	45
4.6	Tikhonov-like Regularization	46
4.6.1	Concept	46

5	Forward Measurement Operators	47
5.1	Linear Polychromatic Approximation	48
5.1.1	Principles	48
5.1.2	Application	48
5.2	Nonlinear Compton and Polychromatic Approximation	51
5.2.1	General problem statement	51
5.2.2	Working results	51
5.2.3	Model descriptions	51
5.2.4	Algorithm description	54
5.2.5	The KN-scattering probability distribution	55
5.2.6	Further Developments	57
5.2.7	Usage and Examples	57
5.3	Geometry and Harmonic Analysis	64
5.3.1	Concept	64
5.4	Fast Ray Casting	66
5.4.1	Concept	66
5.5	Anomalous Diffusion	67
5.5.1	Concept	67
6	Quantitative Object Reconstruction	68
6.1	Mixed-Variable Formulation	69
6.1.1	Principles	69
6.1.2	Application	70
6.2	Linear Polychromatic Reconstructions	71
6.2.1	Results	71
6.3	Reconstructions From Scattering-Effects Corrected Data	83
6.3.1	Results	83
6.4	Nonlinear Polychromatic Reconstructions	85
6.4.1	Results	85

1 Executive Summary

1.1 Motivation

The project named the “X-ray Reconstruction Project” is devoted to the quantitative reconstruction of cylindrically-symmetric objects from single X-ray radiographs under severe limiting constraints on time and computational power. The general problems of interest require material identification and geometric details within several minutes using an off-the-shelf computer. The problem is difficult on several levels including uncertain X-ray source characterization, uncertain detector response and output characteristics, the nonlinear interaction physics describing the X-ray beam evolution, and the need to interpret complex scenes. Rapidly unraveling a detailed object description from a single poorly-understood radiograph is an involved process, to say the least. The task draws upon DDMA’s experience and understanding of such diverse fields as geometric measure theory, radiography, functional analysis, inverse problems, computational science, algorithmics, optimization, statistics, numerical analysis, and, not least, image analysis.

The extended DDMA team gathered together for a directed three-day applied mathematics workshop January 29–31 in Oak Park, Illinois. Our focus was on the assessment of, and solutions to, several specific key problems we feel are necessary to effectively reach our project goal. These three days and the subsequent three week followup work have resulted in significant advances in all of the key areas studied. This introduction continues with synopses of the five study areas. The remainder of this document provides details for many subprojects on mathematical foundations, computational and algorithmic considerations, application examples, and value assessments. This report is not to be interpreted as a project plan or status report; rather, it is a compilation of collective thought and work from a single DDMA workshop.

1.2 Workshop Results

Object Identification Within Radiographs

A key problem in the analysis stream is the identification of objects of interest (OoI) in radiographs and, if possible, to determine a 3-dimensional location and orientation. This ability is central to the accurate interpretation of cluttered scenes.

The map-seeking circuit (MSC) algorithm was applied to this transformation discovery problem. A simple example involving translations and rotations in 2-dimensions was performed on simulated noisy data. These preliminary results suggest that rapid 3-dimensional placement of known objects from a single 2-dimensional radiograph is very likely to succeed. Extensions to identification of object classes is under consideration.

Object Radiograph Extraction

The radiograph of an object of interest (OoI) is often complicated by the presence of other objects in the scene which produce both an additional scattered photon signature and partial

or full occlusions. It is often necessary to extract a radiographic signature that would be present in the absence of these additional objects.

A method of symmetry-enforced total variation inpainting was developed to correct for occlusions. The necessary assumptions are that the OoI has known spherical or cylindrical symmetry and its location and orientation are precisely known. An application to a simulated complex test scene was quite promising: A radiograph of a spherical object was extracted with pleasing accuracy even though a significant portion of the original object image is occluded by a second unknown object.

Scattering Effects Removal

It may also be important to have the capability of removing the photon scattering signature in radiographs. If this can be done successfully then it becomes possible to perform object reconstructions based on linear projection operators. Several methods were developed or proposed. These methods fall into one of two general classes: function interpolation and deconvolution. For these methods to be successful we must not assume any prior knowledge about the objects in the scene and work completely from information in the radiograph signature outside of any objects.

All three developed methods show promise. One method applies a total variation regularized blind deconvolution. By construction it can perform well near object edges, but has difficulty with regions of smooth variation and scenes with objects of different materials. Areas for improvement have been considered. A second method applies heat-kernel smoothing to estimate the scattering in known regions followed by inpainting into unknown regions. The reduction in scattering effects is significant even for complex scenes. A third method, applicable to 1-dimensional data from single symmetric objects, is an interpolation method that assumes a simple Gaussian scattering contribution. It works very well under the assumptions.

Three lines of research were also proposed and are in different stages of development. First, the Gaussian contribution idea is being developed for 2-dimensions and complex scenes. Second, the scattering signature in unknown regions can be obtained by biharmonic inpainting since the boundary conditions on the unknown region are well known. Third, we can use the idea that the scattering signature is a generally low spatial frequency construct and apply a functional optimization technique with a regularization that discourages high spatial frequency solutions.

Forward Measurement Operators

It is important to develop fast forward measurement operators that handle linear and nonlinear effects of beam polychromatics and photon scattering. Our methods must be extremely fast since they will be used within optimization computations. Therefore, our methods will be very approximate but we must strive to incorporate some measure of essential physics.

A simple linear polychromatic operator was developed and it compares favorably to lengthy simulations that ignore nonlinear effects. This operator was effectively used in

optimization-based object reconstruction. Our first true nonlinear forward measurement operator was also developed and tested. Because of computational time limitations it only accounts for singly scattered photons. But it accurately applies scattering and attenuation probabilities and scattering angle distributions. The scattering signature lacks the overall smoothness and low spatial frequency content seen in simulations, but the resultant radiographs are a significant improvement over those produced by linear operators.

Quantitative Object Reconstruction

The measure of success in our research is the ability to produce quantitative object reconstructions, both in material identification and material boundary location. We present here several such reconstructions provided by a mixed-variable optimization (MVO) scheme. Several reconstructions of a single object were performed as a test of new scattering effects correction methods. Both methods corrected for the major deficiency in using raw data – material misidentification. The nonlinear forward Compton measurement operator was also tested, and a reconstruction of the same test object from raw data shows improvement in edge location but difficulties remain for accurate material identification.

1.3 Discussion

This workshop effort has produced a large quantity of useful and promising work of direct applicability and importance to improved quantitative object reconstructions from radiographs. The high quality of the theoretical and algorithmic work, accompanied by standards-driven code is already providing significantly improved applied results and it will likely continue to do so throughout FY06 and beyond.

2 Object Identification Within Radiographs

The goal of this subproject was to develop and demonstrate methods for object-of-interest (OoI) identification and placement from radiographic data and OoI templates. Method requirements include strict computational time constraints (seconds), robustness to OoI occlusion, rotations, 3D translation, scaling, noise, and nonlinear imaging effects.

2.1 Transformation Discovery via Map-Seeking Circuit

Curt Vogel

2.1.1 Principles

The map seeking circuit (MSC) algorithm is a solution to the discrete linear transformation discovery problem. Given some input y , and some output x , we ask the question: what transformation T best satisfies our expectation $T(y) = x$. We consider transformations of the form $T = T_{i_L}^{(L)} \circ \dots \circ T_{i_2}^{(2)} \circ T_{i_1}^{(1)}$, where each $T_{i_\ell}^\ell, i_\ell = 1, \dots, n_\ell$ is linear. The best transformation, say $i_\ell = \hat{i}_\ell$, can be discovered by brute-force exhaustive combinatorial search, but this is not practical as the number of transformations ($n_1 \times \dots \times n_L$) can be quite large. The MSC algorithm can reduce the computational complexity to $\sum n_\ell$.

A detailed discussion on MSC theory can be found in [1]. The flavor of the procedure is the following. MSC computes a comparison measure between a reference and a superposition of transformed input states. Instead of applying a particular discrete transformation at each layer $T_{i_\ell}^{(\ell)}(x)$, MSC considers the superposition of states $\sum_{i_\ell} c_{i_\ell}^\ell T_{i_\ell}^{(\ell)}(x)$ at each layer. Then problem is now formulated as the task of determining the best set of unit vectors $\{c^\ell\}$ that satisfy the problem. This is a much simpler problem to solve because the transformation composition naturally provides an adjoint structure that can be used for gradient ascent optimization on $\{c^\ell\}$.

2.1.2 Application

As an initial test of the MSC algorithm applied to situations representative of object location within radiographs, I considered the task of identifying simple geometric objects (e.g., balls and rectangles) in a cluttered, noisy 2-dimensional image. Prior to this workshop, David Arathorn, Curt Vogel and collaborators at Montana State and UC Berkeley had implemented the MSC algorithm to perform retinal motion tracking [7]. This work required two MSC layers—one for vertical translation and one for horizontal translation to map a given reference image to a target image. During the workshop, I added a third MSC layer to handle rotations to my existing MATLAB codes and I demonstrated that MSC could extract from a cluttered 2-dimensional image a specified ellipse with arbitrary location and arbitrary (angular) orientation. The new codes need to be tested on data provided by DDMA and incorporated into a larger scheme that includes upstream preprocessing (deconvolution and denoising, edge detection, etc.) and downstream postprocessing.

2.1.3 Example

Suppose a processed radiograph takes on the form of the noisy binary image shown in Figure 1. Consider the task of identifying the position (translation and rotation) of the reference object shown in Figure 2. This example considers $n_1 = 88$ vertical translations, $n_2 = 100$ horizontal translations, and $n_3 = 16$ rotations. An exhaustive search of this




Figure 1: Noisy sparsified radiograph containing an object of interest: the small tilted ellipse at the center.

parameter space would require $n_1 n_2 n_3$ comparison calculations. MSC solves this problem in about $10(n_1 + n_2 + n_3)$ comparison calculations. The iteration convergence to the correct transformation vectors $\{c^\ell\}$ is shown in Figure 3. The images show the superposition transformations at each iteration weighted by the current iterate of $\{c^\ell\}$. After about 28 iterations $\{c^\ell\}$ has converged to unit vectors within an acceptable tolerance, and the problem is solved. Known methods for improving the convergence exist but have not been applied in this example. It is reasonable to expect convergence in fewer iterations by about a factor of 4.




Figure 2: The reference description of the object of interest.

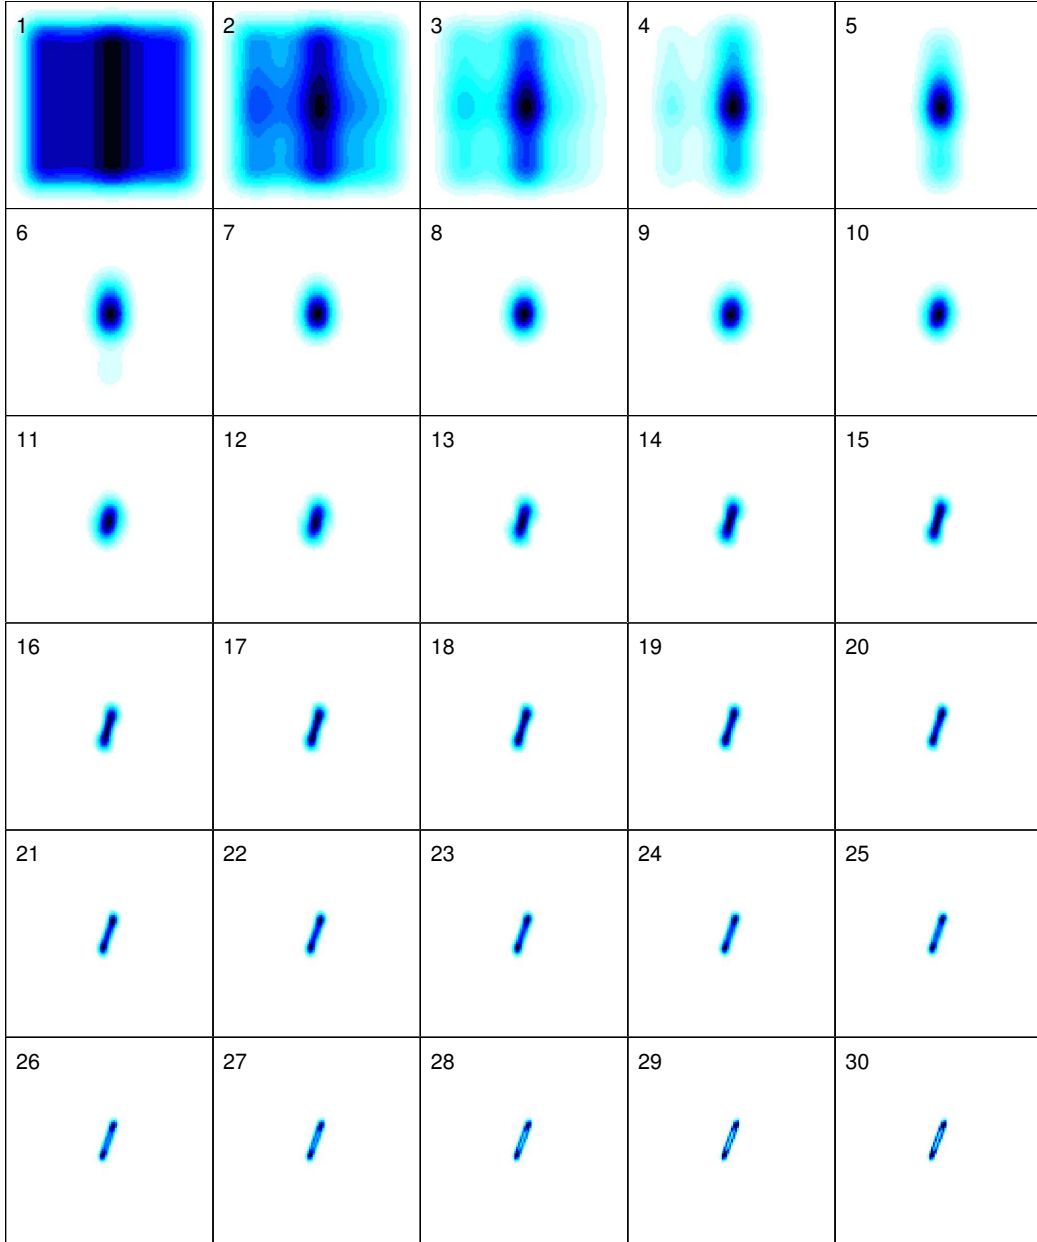


Figure 3: Convergence of the MSC transformation superposition.

3 Object Radiograph Extraction

The goal of this subproject was to develop and demonstrate methods for recovering a radiographic signature of a single object of interest (OoI) from a radiograph of a complex multiple-object scene. The OoI is known to have spherical or cylindrical symmetry, and its exact location and orientation are known. No assumptions about the composition or symmetry of any remaining objects are allowed. Viable algorithms must complete their analysis in less than one minute. Recovery of radiographs both with and without scattering effects are important cases to consider.

3.1 Symmetry-Constrained Inpainting

Selim Esedoglu, Erik Boltt, Bill Allard, Robert Owczarek, Curt Vogel

3.1.1 Principles

We took several approaches to this problem; the one we settled on that gives extremely promising results is based on the following two essential ideas:

1. To figure out the contribution of a given symmetric object to the x-ray radiograph of the overall scene, we propose to *exploit the symmetry* of the object. The symmetry of the object implies that not all functions can be its contribution to the image; we propose using this information in order to narrow down our search space.
2. We remove the object from the scene and *inpaint* the space occupied by it by using information from its exterior, so as to propagate information from other objects in the scene to fill in the region occupied by the given object. The idea is that the inpainted region would now have contributions from other objects in the scene, which can then be subtracted from the original image in order to isolate the contribution of the given object.

Our approach draws in an essential way on ideas from digital image inpainting. Inpainting refers to the task of repairing damaged regions, such as scratches or holes, in images where information is missing. Inpainting repairs the image in these regions by propagating information from surrounding areas where the image is intact. Basically, it is a form of interpolation. However, it differs from standard smooth interpolation in that image data routinely have discontinuities (corresponding to boundaries of objects present in the depicted scene), and these need to be extended appropriately into the inpainting domain by the algorithm.

A well-known and successful image inpainting technique we started with, but then had to eventually significantly modify, is the total total variation based inpainting model of Chan and Shen. It poses the inpainting procedure as finding the minimizer of the following energy:

$$\int_D |\nabla u| + \lambda \int_{D \setminus \Omega} (f - u)^2 dx$$

Here, D is the image domain (which is usually a rectangle representing the computer screen), and Ω is the inpainting domain (such as the space occupied by a scratch) where image information might not be available or is damaged. The given image is denoted by f , which might not even be defined on Ω . Under certain geometric conditions that involve the shape of the inpainting region Ω and the contour lines of f in a neighborhood of Ω , this model successfully propagates f and along with its discontinuities. The parameter λ is to be chosen as large as possible, as the term that it multiplies in the energy is responsible for ensuring that the minimizer agrees with the given image outside the inpainting region.

Experiments with the standard total variation based inpainting model described above quickly indicated that this model is not sufficient for our purposes. The model is quite good when the inpainting region Ω is, for example, long and thin (i.e. has high eccentricity). It is therefore very well suited for instance to scratch removal. Our application, however, calls for reconstructing images in much larger regions than scratches and with arbitrary shapes (that can very well be circular). In such settings, the model cannot propagate the sharp discontinuities of the image f from outside Ω to inside. An example of this is given below. This caveat makes it crucial to incorporate additional information about the image to be reconstructed into the model.

Let f denote the areal attenuation image corresponding to the original scene with the object to be removed in it. Let u denote the areal attenuation image corresponding to the scene without the object to be removed; this is the unknown to be found by our proposed algorithm. Also, let Ω be the space occupied by the object to be removed. Our basic approach is to obtain u by treating Ω as an inpainting region and propagating f from outside of Ω to inside. We accomplish this via a new total variation based inpainting model. The resulting image is our approximation to u . Once u is obtained this way, $f - u$ constitutes our approximation to the attenuation image corresponding to a scene with only the object to be removed in it.

An important piece of information is that the density of the object to be removed from the scene has a certain symmetry. The new image information generated by the inpainting algorithm in this space must be compatible with the known symmetry of the removed object. For example, if we assume that the object to be removed from the scene has the shape of a sphere and spherically symmetric density, then its contribution to areal attenuation image will also have radial symmetry. This implies that $f - u$ must have radial symmetry inside Ω . A novel aspect of our approach to inpainting is to incorporate this type of information into the model.

We illustrate our approach by considering in detail the situation the object to be removed occupies a ball in the image domain and makes a circularly symmetric contribution to the areal attenuation image. In that case, the inpainting domain Ω is given by a ball $B_R(0)$ of radius R , which we take to be centered at the origin. Our modified total variation based inpainting model in this case consists of finding the minimizer of the following energy:

$$\begin{aligned} \int_D |\nabla u| \, dx + \lambda \int_{D \setminus B_R(0)} (u - f)^2 \, dx \\ + \mu \int_{B_R(0)} \left((u - f) - \frac{1}{|\partial B_{|x|}(0)|} \int_{\partial B_{|x|}(0)} (u - f) \, d\sigma \right)^2 \, dx \end{aligned}$$

The third term in the energy constitutes its essential difference from the standard total variation inpainting model. It encourages the minimizer u of the energy to be such that its difference $u - f$ from the original image f should have circular symmetry in $B_R(0)$. The third term encourages circular symmetry by penalizing deviation of $u - f$ at any point

$x \in B_R(0)$ from its average calculated on the circle centered at the origin on which x lies. In the formula above, $\partial B_{|x|}(0)$ denotes the circle of radius $|x|$ centered at the origin, and $|\partial B_{|x|}(0)|$ denotes its length. The parameter μ is to be chosen by the user; it determines the strength of the symmetry term in relation with others in the energy. Gradient descent for this energy reads:

$$u_t = \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right) + \lambda \mathbf{1}_{D \setminus B_R(0)}(x)(f - u) \\ + \mu \mathbf{1}_{B_R(0)}(x) \left(\frac{1}{|\partial B_{|x|}(0)|} \int_{\partial B_{|x|}(0)} (u - f) d\sigma - (u - f) \right)$$

We approximate the minimizer of the energy by solving this equation for large times, starting from an appropriate initial guess.

A simplified version of the variational problem above has the form

$$\int_D |\nabla u|^2 dx + \lambda \int_{D \setminus B_R(0)} (u - f)^2 dx \\ + \mu \int_{B_R(0)} \left((u - f) - \frac{1}{|\partial B_{|x|}(0)|} \int_{\partial B_{|x|}(0)} (u - f) d\sigma \right)^2 dx$$

where the total variation term $\int |\nabla u|$ is replaced with the more classical Dirichlet energy $\int |\nabla u|^2 dx$. This model does not propagate sharp edges into the inpainting region. Its advantage is that it is very fast to compute. Hence it can be used to generate a good initial guess for the total variation based model described above.

One of the biggest technical challenges involved in the numerical treatment of the models above is implementing the spherical means procedure that appear in them in a numerically convenient yet accurate manner. Beyond that, Euler-Lagrange equation for the new total variation based model is a nonlinear, degenerate elliptic PDE. For efficient solution of the gradient descent scheme, a linear elliptic solver is required using MATLAB's backslash function. This has been implemented and will be incorporated into the gradient flow code subsequently.

3.1.2 Generalized Symmetry Discussion

Even if an object has a regular geometric shape such as a cylinder, a cone, or a sphere, and density with the corresponding symmetry, its radiographic projection need not have any simple symmetry. Most orientations in space would destroy the symmetry. For example, if the object is a cylinder with cylindrically symmetric density, in case it happens not to be oriented in a perfectly normal manner to the image plane, its projection will no longer have cylindrical symmetry. However, the class of functions that describe such projections is still very limited. Therefore, building this information into the inpainting model as we

propose would still lead to significant improvement, as it would prevent one from looking for the solution among all possible functions; instead, one would look for it among the class of functions compatible with projections of symmetric densities, which is a much smaller class. Of course, the term to be added to the energy for this purpose would no longer be as simple as the circular means based term that we added. However, there is a simple way to do it for any shaped object with density of some symmetry, oriented in any way in three space.

3.1.3 Example

We now show a sample computation obtained using the new total variation based inpainting method proposed above. The given radiography intensity image is shown in Figure 4(a). It depicts a scene with several objects in it. The first step is to convert to an areal attenuation image by taking logarithms. The resulting image, which we will denote f in what follows, is displayed in Figure 4(b). The region Ω to be inpainted by our model will be taken to be the space occupied by the circular object nearest the center of the image.




Figure 4: The given radiography intensity image (a) and the corresponding areal attenuation image (b).

If we apply the standard total variation based image inpainting algorithm to the areal attenuation image, the result is quite unacceptable, as is shown in Figure 5.

As can be seen in Figure 5(a), the bar is not propagated into the circular region at all. This is a natural and well-known drawback of total variation based inpainting: Propagating the edges of the bar far into the circular region would incur too high a cost. Although there are more complicated inpainting models that try to address this issue (for instance by incorporating higher order terms in the energy), they would all have trouble with as challenging a situation as this one, and are computationally a nightmare. When this un-

satisfactory completion is subtracted from the original attenuation image, so as to isolate the contribution to attenuation from the object in the circular region, the resulting image is shown in Figure 5(b). We see that because standard total variation based inpainting failed to propagate the bar into the circular region, it shows up as if it were part of the contribution of the circular object to the radiograph. Figures 5(c) and 5(d) show the intensity images corresponding to the unsatisfactory attenuation images computed by the standard total variation inpainting model.

We now turn to the results obtained by the model we developed, which incorporates the circular means terms into total variation inpainting model in order to encourage symmetry where it is known to hold. The inpainting of the outside objects obtained by minimizing our proposed model is shown in Figure 6. As can be seen in Figure 6(a), the bar begins to be propagated into the circular region, despite its aspect ratio that is so unfavorable from the point of view of total variation. More spectacularly, Figure 6 shows the result when this reconstruction is subtracted from the given attenuation image. The result is a nearly circularly symmetric areal attenuation.

Figures 6(c) and 6(d) show the intensity images corresponding to the attenuation images obtained by our proposed model that are shown in Figures 6(a) and 6(b). Note the significant improvement over Figure 5 and that were obtained using the standard total variation based model.

3.1.4 Implementation Details

Here we give details of how the computational examples shown were obtained. First, the intensity image g is converted to areal attenuation, which then denoted f :

$$f = \log(g + 0.001);$$

Then, we normalize the image f so that it is between 0 and 1, making note of the transformation:

$$\begin{aligned} m &= \min(\min(f)); \\ f &= f - m; \\ a &= \max(\max(f)); \\ f &= f/a; \end{aligned}$$

The inpainting domain will be the circular region occupied by the object nearest the center of the image shown in Figure 4(a). The full image (not shown here) has a resolution of 300x300, and the circular region in question has radius 30 pixels and is centered in the original image at 91,211. We therefore now prepare a matrix R that indicates the complement of this region (mathematically, it will approximate the term $\mathbf{1}_{D \setminus \Omega}(x)$ in the proposed model):

$$\begin{aligned} [i,j] &= \text{meshgrid}(1:300); \\ R &= (i-91).^2 + (j-211).^2 >= 30^2; \end{aligned}$$

We can now call the routine `inptvsym.m` to carry out inpainting via the proposed model:

```
u = inptvsym(3000,1/(5*300^2),f,f,R,91,211,30,10000,6000);
```

Here, the first argument of 3000 is the number of time steps to take. The second argument, $1/(5 * 300^2)$, is the time step size. In the current implementation, the gradient descent equation is solved via explicit time stepping, which puts the restriction that $\delta t < \frac{\delta x^2}{4}$. Once the implicit elliptic solver gets incorporated using lagged diffusivity, this restriction will be lifted, allowing for much faster computation of the stationary state. The third argument is the initial condition. Since the original image f is a reasonable place to start, we pass in that in this slot. The fourth argument is the original image. The sixth and seventh arguments are the coordinates of the center of the circular region, while the ninth argument is its radius. The final two arguments are the parameters λ and μ , in that order. The output variable `u` now contains the approximation to the minimizer. To obtain the corresponding intensity image for `u`, one can form the image:

```
exp(-(au+m))-0.001
```




Figure 5: The solution found by the standard total variation inpainting algorithm used on the areal attenuation image of Figure 4(b). The top figures show the areal attenuation of the background (a) and object of interest (b). The bottom figures are the corresponding intensity images.




Figure 6: The solution found by the new total variation inpainting based algorithm we propose, used on the areal attenuation image of Figure 4(b). The top figures show the areal attenuation of the background (a) and object of interest (b). The bottom figures are the corresponding intensity images. The result shows a marked improvement over that obtained by the standard model; compare with Figure 5.

4 Scattering Effects Removal

The goal of this subproject was to develop and demonstrate methods for identifying and removing the scattering effects signal in radiographs. The guiding assumptions are to assume no prior knowledge about the composition or symmetry of objects in the scene, and the absolute background response in the absence of photon scattering is known. Viable algorithms must complete their analysis of potentially complex scenes in less than one minute.

4.1 Total Variation Based Deconvolution

Paul Rodriguez, Ousseini Lankoande

4.1.1 TV Overview

Let u be a function defined in $[0, 1]$, then its total variation is defined by $TV(u) = \sup \sum_k |u(x_k) - u(x_{k-1})|$, $x_k \in [0, 1]$ and the supremum is taken over all possible partitions in $[0, 1]$. In a discrete setup, we use the approximation

$$TV(\mathbf{u}) \approx \sum_{k=1}^N |u_k - u_{k-1}| \quad (1)$$

where $\mathbb{R}^1 \ni \mathbf{u} = [u_k]$, $k = 0, 1, \dots, N$. Then the minimization of the TV-functional is given by

$$T_\alpha(\mathbf{u}) = \|K\mathbf{u} - \mathbf{d}\|_2 + \alpha TV(\mathbf{u}) \quad (2)$$

where \mathbf{u} is solution to the minimization problem, K is a matrix, and \mathbf{d} is the observed (noisy) data. In order to overcome the non-differentiability of the Euclidean norm (eq. 1) at the origin we use the smooth approximation $\Upsilon_\epsilon(u) = \begin{cases} \frac{u}{2\sqrt{\epsilon}} & u \leq \epsilon^2 \\ \sqrt{u - \epsilon} & u > \epsilon^2 \end{cases}$; then eq. 2 is transformed into

$$T_\alpha(\mathbf{u}) = \|K\mathbf{u} - \mathbf{d}\|_2 + \alpha J_\epsilon(\mathbf{u}) \quad (3)$$

where $J_\epsilon(\mathbf{u}) = \sum_{k=1}^N \Upsilon_\epsilon((D_k \mathbf{u})^2)$ and $D_k \mathbf{u} = u_k - u_{k-1}$. Note that eq. 3 can be easily minimized (since J is Frchet differentiable i.e.: $\text{grad} J(\mathbf{u})$ can be computed by evaluating $\left. \frac{\partial J(\mathbf{u} + \tau \cdot \mathbf{v})}{\partial \tau} \right|_{\tau=0} = \langle \text{grad} J(\mathbf{u}), \mathbf{v} \rangle$) using any optimization technique (such as steepest descent, Newton's method, or lagged diffusivity fixed point iteration, etc.)

The two dimensional case can be easily extended. First consider that matrix K is the discretization of a linear operator which acts on functions of two variables. Also, let \mathbf{u} be the lexicographical column ordering of the two-dimensional data set $[u_{n,k}]$, where $n \in [0, N_y]$ and $k \in [0, N_x]$. Then the discrete two-dimensional smooth approximation of the TV can be written as

$$J_\epsilon(\mathbf{u}) = \sum_{n=1}^{N_y} \sum_{k=1}^{N_x} \Upsilon_\epsilon \left((D_{n,k}^x \mathbf{u})^2 + (D_{n,k}^y \mathbf{u})^2 \right) \quad (4)$$

where

$$D_{n,k}^x \mathbf{u} = u_{n,k} - u_{n,k-1} \quad D_{n,k}^y \mathbf{u} = u_{n,k} - u_{n-1,k}. \quad (5)$$

Then the TV-functional can be written exactly (with abuse of notation) as eq. 3.

In order to solve the minimization problem, we need to compute the gradient of $J_\epsilon(\mathbf{u})$ (note that $\text{grad} T(\mathbf{u}) = K^T \cdot (K\mathbf{u} - \mathbf{d}) + \alpha \cdot \text{grad} J(\mathbf{u})$). For the two-dimensional case, it

can be easily shown that:

$$\text{grad}J(\mathbf{u}) = L(\mathbf{u})\mathbf{u} \quad (6)$$

$$L(\mathbf{u}) = \left(D_x^T \text{diag}(\Upsilon'_\epsilon(\mathbf{u})) D_x + D_y^T \text{diag}(\Upsilon'_\epsilon(\mathbf{u})) D_y \right) \quad (7)$$

where D_x and D_y are matrices corresponding to the grid operations in 5, and $\text{diag}(\Upsilon'_\epsilon(\mathbf{u}))$ denotes the diagonal matrix with entries $\Upsilon'_\epsilon \left((D_{n,k}^x \mathbf{u})^2 + (D_{n,k}^y \mathbf{u})^2 \right)$.

In [6] there is a nice compilation of several algorithms (described using matrix notation) to solve eq. 3 for one/two dimensional cases. In particular, the lagged-diffusivity algorithm is explicitly described.

Algorithm 1 Lagged diffusivity fixed point

\mathbf{u}_0 : initial guess

for $n = 0, 1, \dots$

$L_n = L(\mathbf{u}_n)$

$g_n = K^T(K\mathbf{u}_n - \mathbf{d}) + \alpha L_n \mathbf{u}_n$

$H = K^T K + \alpha L_n$

$\Delta_n = -H^{-1} g_n$

$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta_n$

4.1.2 Method

Total variation has been successfully applied for blind deconvolution problems [2]. Here we propose a modification to the TV functional (eq. 3) which will lead to a deconvolution problem.

Physical model for scattered radiation in X ray imaging is fairly complex. Here we do not attempt to describe known models nor to propose a novel one. What follows is based on an effect that can be observed in (digital) X ray images: scatter is more visible at the edges of the objects (e.g. edges are diffused). Considering the previous statement, we can simulate (in a very crude fashion) the effect of scattering given a digital (scatter-free) image by adding to it the result of the convolution of a Gaussian filter (properly scaled) with the absolute value of the gradient of the original image.

Figure 7 shows the result of a simulation that implements the idea expressed in the previous paragraph. Note that the (temporal/spatial) extension of the scatter will depend on the number of elements of the Gaussian filter and its intensity will depend on the l_1 norm of the Gaussian filter.

4.1.3 Total Variation Descattering

Based on the proposed forward model (presented in the previous section), we propose to modify TV functional (see eq. 3) to include, in the fidelity term, the effect of a Gaussian





Figure 7: Scatter simulation. Let x be the scatter free signal (in blue) and z be the (simulated) scattered signal, then $z = \text{conv}(h, |\nabla x|) + 0.001 \cdot \eta$, where h is a Gaussian filter (as described in fig. 8) and η is zero mean Gaussian noise



Gaussian filter

Figure 8: Gaussian filter used for TV de-scattering. Filter h was generated $h = \exp(-0.5 * (-2 : 0.05 : 2)^2)$ and then scaled so that $\|h\|_1 = 1.1$.

filter that diffuses the edges of the input image. Eq. 8 is the propose TV de-scatter cost function:

$$C_\alpha(\mathbf{u}) = \left\| \mathbf{u} + H * \sqrt{(D_{n,k}^x \mathbf{u})^2 + (D_{n,k}^y \mathbf{u})^2} - \mathbf{d} \right\|_2 + \alpha J_\epsilon(\mathbf{u}) \quad (8)$$

where $*$ denotes convolution, H is a 2D separable filter (assumed to be Gaussian) e.g. H is equal to the tensor product of h with itself, where $h = a \cdot \exp(-0.5 \cdot \nu)$, a is some constant, and ν is a vector whose values sample the real line. Even though eq. 8 exactly implements the forward model (described in the previous section), note that $|\nabla \mathbf{u}| = \sqrt{(D_{n,k}^x \mathbf{u})^2 + (D_{n,k}^y \mathbf{u})^2}$ is a non-invertible, non-linear function and therefore we can not expressed eq. 8 as eq. 3. Also note that matrices D_x and D_y (corresponding to the grid operations defined in 5) are non-invertible matrices; for the fidelity term we propose the following grid operations:

$$\check{D}_{n,k}^x = 0.5 \cdot u_{n,k-1} - 2.0 \cdot u_{n,k} + 1.5 \cdot u_{n,k+1} \quad (9)$$

$$\check{D}_{n,k}^y = 0.5 \cdot u_{n-1,k} - 2.0 \cdot u_{n,k} + 1.5 \cdot u_{n+1,k} \quad (10)$$

there are many choices to approximate the first derivative by finite differences; we chose 9 and 10 because of their simplicity and stability. A well-known approximation, often mentioned in the TV literature but used for other purposes (i.e. to approximate $TV(\mathbf{u})$) is $|\nabla \mathbf{u}| \approx |D_{n,k}^x \mathbf{u}| + |D_{n,k}^y \mathbf{u}|$; inspired by this here we propose $|\nabla \mathbf{u}| \approx S^x \check{D}_{n,k}^x \mathbf{u} + S^y \check{D}_{n,k}^y \mathbf{u}$ where S^x and S^y are matrices such that $S^x \check{D}_{n,k}^x \mathbf{u} = |\check{D}_{n,k}^x \mathbf{u}|$ and $S^y \check{D}_{n,k}^y \mathbf{u} = |\check{D}_{n,k}^y \mathbf{u}|$; then, we modified eq. 8 so that the TV de-scattering cost functional is defined as:

$$C_\alpha(\mathbf{u}) = \left\| \mathbf{u} + H * \left(S^x \check{D}_{n,k}^x \mathbf{u} + S^y \check{D}_{n,k}^y \mathbf{u} \right) - \mathbf{d} \right\|_2 + \alpha J_\epsilon(\mathbf{u}) \quad (11)$$


we must stress that matrices S^x and S^y are not fixed (they depend on the actual values of \mathbf{u}); nevertheless, from a practical point of view we will *think* of them as fixed, and express eq. 11 as:

$$C_\alpha(\mathbf{u}) = \left\| \left(I + H_M S^x \check{D}_{n,k}^x + H_M S^y \check{D}_{n,k}^y \right) \mathbf{u} - \mathbf{d} \right\|_2 + \alpha J_\epsilon(\mathbf{u}) \quad (12)$$


where H_M is a matrix gives identical results to the convolution operation. Eq. 12 can be solved using any well-known algorithm that solves eq. 3, in particular the proposed TV de-scatter cost functional (eq. 12) was implemented in C and made compatible with IDA, using the lagged-diffusivity fixed-point algorithm (see algorithm 1).

4.1.4 Results


Here we present the simulation results obtain by the proposed TV de-scatter method. The method was applied to scatter images generated by MCNP which doesn't follow the proposed forward method, but rather approximates an actual experiment by particle transport calculations.



(a) Scene1-10 (scatter free)



(b) Scene1-10 (scatter)




(c) Scene1-10 (TV de-scatter)


Figure 9: Results for TV de-scattering.

Figure 9.a shows a scatter-free image which contains several solid objects of different materials. Figure 9.b is its scatter version; figure 9.c is the result of the TV de-scatter method (with parameters $\epsilon = 0.01$, $\alpha = 0.01$, filter h as describe in figure 8). More interesting, in figure 10 the profiles for all three versions are shown for three different rows; it can be observed that even though results are not prefect, the proposed method behaves as expected: it denoises and reduces the scatter of the noisy image. The input image (figure 9.b) size is 300×300 , to complete 5 iterations (in order to obtain the results) takes about 3 seconds on a Pentium4 running at 3GHz, with 1G of RAM.


Figure 11.a is pretty similar to 9.a but it contains a mixture of solid and hollow objects; materials are the same (as in figure 9.a). Figure 11.b is its scatter version; figure 11.c is the result of the TV de-scatter method (with parameters $\epsilon = 0.01$, $\alpha = 0.01$, filter h as describe in figure 8). In figure 12 the profiles for all three versions are shown for three different rows



(a) Profile (row 80)



(b) Profile (row 150)




(c) Profile (row 200)


Figure 10: Profiles for the images shown in fig. 9.

(compared it with fig. 10). The input image (figure 11.b) size is 300×300 , to complete 5 iterations (in order to obtain the results) takes about 3 seconds on a Pentium4 running at 3GHz, with 1G of RAM.


Figure 13.a shows a different scene than the two previous examples; the most remarkable object is the step wedge (right of the image). Figure 13.b is its scatter version; figure 13.c is the result of the TV de-scatter method (with parameters $\epsilon = 0.01$, $\alpha = 0.1$, filter h as describe in figure 8) and figure 13.d is the result of the TV de-scatter method (with parameters $\epsilon = 0.01$, $\alpha = 0.05$, same filter h). In figures 14 and 15 show the profiles for both results. Note the effect of parameter α (the smaller the value of α the less big jumps are penalized) when both results are compared. The input image (figure 13.b) size is 1024×1250 , to complete 5 iterations (in order to obtain the results) takes about 7 minutes (when $\epsilon = 0.01$, $\alpha = 0.1$) and 6.3 minutes ($\epsilon = 0.01$, $\alpha = 0.05$) on a Pentium4 running at



(a) Scene2-10 (scatter free)



(b) Scene2-10 (scatter)




(c) Scene2-10 (TV de-scatter)

Figure 11: Results for TV de-scattering.


3GHz, with 1G of RAM.

4.1.5 Conclusions


- The approximation $|\nabla \mathbf{u}| \approx S^x \check{D}_{n,k}^x \mathbf{u} + S^y \check{D}_{n,k}^y \mathbf{u}$ is justified from a practical point of view.
- Results depend of the parameters of the Gaussian filter $h = a \cdot \exp(-0.5 \cdot \nu)$: a controls the gain of the scatter and the number (and location) of samples ν controls the spread of the scatter.
- For more accurate results, we propose to use a spatial varying Gaussian filter whose parameters depend of the physical properties of the particular material being exposed



(a) Profile (row 80)




(b) Profile (row 150)




(c) Profile (row 200)

Figure 12: Profiles for the images shown in fig. 11.


to the X ray flux.




(a) Scene2-10 (scatter free)



(b) Scene2-10 (scatter)




(c) Scene2-10 (TV de-scatter $\alpha = 0.1$)




(d) Scene2-10 (TV de-scatter $\alpha = 0.05$)


Figure 13: Results for TV de-scattering.



(a) Profile (row 800)




(b) Profile (column 700)




(c) Profile (column 1100)


Figure 14: Profiles for the images shown in fig. 13.a, 13.b and 13.c



(a) Profile (row 800)



(b) Profile (column 700)



(c) Profile (column 1100)

Figure 15: Profiles for the images shown in fig. 13.a, 13.b and 13.d

4.2 Heat-kernel Regularization and Inpainting by Linear Regression

Pete Schultz, Kevin Vixie

4.2.1 Principles

The standard approach to X-ray radiography is to assume that photons either travel straight through the object unimpeded, or are absorbed by the radiographed object and thus fail to reach the image plane. The expected radiographic intensity at an image point x is therefore $f_{NS}(x) = f_0 - f_{abs}(x)$ where f_0 is the expected intensity from radiographing a blank scene, and $f_{abs}(x)$ is the intensity loss due to absorption.

However, this treatment ignores scattering effects. (“NS” is for no scattering.) Some of the intensity loss is due to X-rays being scattered to another point in the image plane, instead of being absorbed. The actual radiographic intensity is therefore expected to be greater than f_{NS} . A better model would be $f_0 - f_{abs}(x) + f_{sc}$, where f_{sc} is the contribution from scattering.

In addition, there is Poisson noise in the real radiograph, so that the model for radiography becomes $f(x) = f_0 - f_{abs}(x) + f_{sc} + f_{noise}$. We were given the task of determining f_{sc} , given $f(x)$ and f_0 . One approach we took was heat-kernel regularization and inpainting by linear regression.

The first idea is that we can separate the regions of the image with objects above it, from regions with no objects above it. Label the region of the image plane with no objects above it, A_{img} , and the corresponding region in the source plane A_{src} .

An X-ray beginning at a point in A_{src} in the source plane proceeds unimpeded to the corresponding point in A_{img} . At points in A_{img} , then, $f_{abs}(x)$ is zero. Photons originating in B_{src} can be scattered off the object to any point in A_{img} or B_{img} , so that $f_{sc} > 0$ at every point in the image plane.

Therefore the expected radiographic intensity (ignoring noise) is greater than f_0 at every point in A_{img} . On the other hand, since absorption effects are expected to be greater than scattering effects, we expect that at most points of B_{img} will have $f < f_0$. To a first approximation, we use the given value f_0 to separate A_{img} and B_{img} . The approximation can be refined to take into account the effect of Poisson noise and areas where $f_{abs}(x)$ are close to zero. The details of this refinement are discussed in the Application section.

Since $f_{abs} = 0$ on A_{img} , we know $f_{sc} + f_{noise}$ on A_{img} .

The second idea is that $f_{sc}(x)$ and its derivatives are smooth. The distribution of X-rays scattering off a single point scatterer will be a smooth function. For a typical scene, the scattering will be an integral over all scatterers, of single-point distributions, and will therefore also be smooth.

This idea has two consequences. First, we can separate f_{sc} and f_{noise} on A_{img} by regularization. In this way, we reconstruct f_{sc} on A_{img} .

The second consequence is that f_{sc} is smooth even across the interface between A_{img} and B_{img} , even though the radiograph can show sharp edges. This means that we can get

a good approximation to f_{sc} in B_{img} by continuing smoothly across the boundary, even without knowing f_{abs} in B_{img} .

This discussion suggests an algorithm for reconstructing f_{sc} from f_0 and $f(x)$:

1. Partitioning: identify A_{img} and B_{img} .
2. Regularization: reconstruct f_{sc} on A_{img} by regularizing $f(x)$. We regularized by convolution with a heat kernel, using only points in A_{img} .
3. Inpainting: reconstruct f_{sc} on B_{img} by smoothly continuing f_{sc} across the interface between A_{img} and B_{img} . To do so, we identify pixels on the boundary of the region where f_{sc} is known; and estimate f_{sc} at these points.

We then update the set of points at which f_{sc} is known, and repeat the inpainting procedure, until all points have been filled.

4.2.2 Application

The MATLAB function m-file `sei.m`¹ calculates the scattering contribution, given a radiograph and the reference level f_0 . It calls the following MATLAB functions.

`maskobjects.m` takes the radiograph and f_0 , and returns a masked radiograph, where the values at points in B_{img} are replaced with zeros.

`planarregression.m` takes a collection of (x, y, z) triplets, and a corresponding weights, and performs a weighted least-squares regression to find the plane that best fits the points.

Usage: the routine is called as follows:

```
u = sei(A,blanklevel,kw,pictures)
```

A is the input radiograph. It is a matrix.

blanklevel is the radiographic intensity of a radiograph of nothing. It is a scalar.

kw = kernel width. Default is 7. When the routine performs a center-weighted average, or a center-weighted linear regression, the kernel width determines the size of the radius to consider. The weighting function is a Gaussian.

pictures = a diagnostic flag. Default is zero. If pictures=1, the routine will draw each step of the inpainting. Otherwise, it won't.

The output u is the estimate of f_{sc} . The estimate of a theoretical scattering-free radiograph, from a given radiograph and blank level, would be

```
NS = SC - sei(SC,blanklevel)
```

The details of the algorithm are as follows.

Identifying A_{img} . The routine `sei.m` calls `maskobject.m` to separate A_{img} and B_{img} . The first approximation to A_{img} is to identify those points at which $f(x) > f_0$. This ignores two issues however.

¹The function `sei` is short for "scattering effects isolation" and should not be confused with any reference to Tom's father, Sei Asaki.

The first is Poisson noise, which means that a point in A_{img} could have $f(x) < f_0$. Due to the nature of Poisson noise, however, such points are expected to be isolated. If a point identified in B_{img} has three or more neighbors in A_{img} , that point is assumed to also belong to A_{img} .

The second is that at the edges of B_{img} , the value of f_{abs} might be less than the scattering contribution f_{sc} . Consequently the expected value of $f(x)$ can be greater than f_0 , leading to false inclusion in A_{img} .

To correct for this issue, the routine identifies the points on the boundary of A_{img} , and looks at nearby points in A_{img} but not on the boundary. It then takes a center-weighted average of these points. If the value of $f(x)$ at the point in question is at least two standard deviations of the Poisson noise below the average of nearby points, then the point is assigned to A_{img} . The standard deviation is estimated numerically from the distribution of points in A_{img} .

The routine `maskobject.m` returns a masked radiograph, where the values in A_{img} are unchanged, and the values in B_{img} are replaced by zeros.

Regularizing $f(x)$ in A_{img} .

The values on A_{img} still include Poisson noise at this point. To correct for this, the values on A_{img} are convolved with a heat kernel. The convolution is done in such a way that points in B_{img} do not effect the regularization. If G is the convolution kernel, the regularized value f_{reg} is given by

$$f_{reg}(x) = \frac{\int_A f(y)G(x-y)dy}{\int_A G(x-y)dy},$$

where \int_A means the integration is over the region A_{img} .

On A_{img} , the scattering contribution is now taken to be $f = f_{reg} - f_0$.

Inpainting f_{sc} in B_{img} . The inpainting routine proceeds in a loop:

1. Identify points on the boundary of the set where f_{sc} is unknown.
2. At each point in the boundary, identify points within `kw` pixels at which f_{sc} is known.
3. Perform a weighted linear regression through these points (using `planarregression.m`), to find the best-fit plane for the values of f_{sc} , and determine the value corresponding to the boundary pixel
4. Assign the appropriate values of f_{sc} to the boundary points.
5. Repeat. The loop continues until every point has a value of f_{sc} assigned to it.

Run times for executing `sei.m` on the 300x300 radiographs in the test data were approximately 45-50 seconds on a 1 GHz Mac PowerBook G4.




Figure 16: Test radiograph

4.2.3 Examples

We present results of the above algorithm applied to test data `Scene1.mat` we were provided in the workshop. This scene is illustrated in Figure 16. For this scene, radiographs were simulated both neglecting and including scattering. The scene is a cube 20cm on a side and the image plane is a square 30cm on a side. The images are 300x300 pixels in size. Results were simulated using several values for the distance from the location to the image plane. We present results from the radiographs `scene1_ns_010` and `scene1_sc_010`, which, placed the image plane 10cm from the center (that is, right up against the edge of the cube). The “sc” designation indicates that scattering was included in the simulation; the “ns” designation indicates that it was neglected.

Figure 18 shows the values of these two test images along the vertical cross section shown in Figure 17. Figure 19 shows the values after subtracting the scattering contribution computed by `sei.m`. The improvement is clear. This cross section was chosen for illustration because it was the vertical cross section with the *largest* root-mean-square distance between




Figure 17: Comparison of simulation with scattering effects excluded (dots) and scattering effects included and then subtracted off (solid).

the no-scattering simulation and the reconstruction.

The table below shows the level of improvement obtained when `sei.m` is applied to the various test radiographs provided. The columns are as follows:

- Column 1: name of radiographic scene (xx is a placeholder for both sc and ns)
- Column 2: RMS distance between simulations including and neglecting scattering.
- Column 3: RMS distance between no-scattering simulation and `sei.m` reconstruction.
- Column 4: Percentage decrease from column 3 to column 4.




Figure 18: Comparison of simulations with scattering effects included (solid line) and excluded (dots).




Figure 19: Comparison of simulation with scattering effects excluded (dots) and scattering effects included and then subtracted off (solid).

Name	NS vs. SC	NS vs. reconstruction	Percentage improvement
scene1_xx_010	3.69E-03	7.19E-04	80.5
scene1_xx_020	1.72E-03	1.54E-04	91.1
scene1_xx_050	6.26E-04	1.43E-04	77.2
scene1_xx_100	3.45E-04	1.60E-04	53.7

4.3 1D Smooth Function Inpainting

Robert Owczarek

4.3.1 Principles

We consider the one-dimensional inpainting or interpolation problem for symmetric scattering signature removal. The crucial observation is that the scattering component of the signal is a continuous function, which, when noise is neglected, is smooth and slowly changing with distance from the center, where it has maximal value. We neglect any additional concepts of the underlying physics. The idea is to reconstruct the function in the region of the object from full knowledge of the scattering signature outside of the object. We assume exact knowledge of the background intensity. By subtracting from an image, where there is both attenuation and scattering, the known blank (limited to the section), we get a signal where outside of the body there is only scattering component, while at the body area there is both scattering and attenuation components. Attenuation is unknown and we want to recover it. Continuation of the scattering function inside the body is our main goal. We initially assume that the scattering signature can be reasonably modeled by a single Gaussian. This assumption can be relaxed in the future, but preliminary results are encouraging. We consider a Gaussian of the form

$$f(x) = A \exp\left(-\frac{x^2}{\sigma^2}\right) \quad (13)$$

with A and σ unknown parameters. We can use our partial knowledge of the function $f(x)$ to recover the values of the parameters. The value of σ can be found by the observation that the points of inflection are as a rule in the region outside the body, where the function is known. The estimate of the points where the second derivative vanishes is rather rough but led to surprisingly good results, so it appears worthy of further study. There is a simple relation between the points of inflection and σ , namely, $\sigma = \pm x_{\pm}$, where x_{\pm} are the two values of x for the points of inflection. Further calculations are not very sensitive to this value. We assume that values of x in the region of the object are small, so that the Gaussian can be approximated as

$$f(x) = A\left(1 - \frac{x^2}{\sigma^2}\right). \quad (14)$$

Since we know σ ; we impose the condition that at the boundary of the body the extrapolation (approximately) continuously matches the external portion of the function. This in principle allows to get the value of A . In practice, noise in $f(x)$ near the boundary complicates continuous function matching. So, the match should be made to some average value, or otherwise established one. The improvement of the reconstruction can be made by non-linear fitting procedure, with initial data supplemented as above. Then the extrapolation of the scattering function is subtracted from the section representing initial image, where all problems were present. The resulting cleaned up section of the image still has some noise, but in principle is free of scattering. The above procedure is not well automated at present and much work remains to be done.




Figure 20: Initial profile of the image with scatter.

4.3.2 Application

This procedure was applied to the image `sph_sc_tr` – a good example with a significant scattering signature and a significant amount of noise. The procedure is sensitive to the extent of the scattering signature that is available outside the object. In particular, it is necessary to have an estimate of the inflection points and a sufficient amount of data for fitting.

4.3.3 Examples

The initial intensity profile is shown in Figure 20. The scattering signature, including the smooth interpolation, is shown in the Figure 21. Subtraction of this scattering function from the initial intensity profile gives an estimate of the attenuation intensity radiograph. Figure 22 shows comparison of all three profiles: attenuation image; scattering-effects corrected image; and image with scattering effects.




Figure 21: Profile of scattering after extrapolation.




Figure 22: Comparison of reconstructions for clean image and cleaned up image.

4.4 n-Gaussians Model of Scattering Signatures

David Caraballo, Kevin Vixie, Paul Rodriguez

4.4.1 Principles

We are trying to remove the scattering signature from a 2D radiograph image. When objects are masked from the image, the remaining scatter signature looks roughly like a sum of bell-shaped, thick-tailed distributions (not quite Gaussian). By accurately modeling this scatter, we should be able to remove much of it.

The main idea is to use data to perform a non-linear, least-squares regression, fitting the scatter data to a sum of n Gaussians, where n is relatively small ($n \sim 5$). This should suffice for a relatively small number of essentially convex objects. If there are a large number of highly non-convex objects, then a large n would be required, and this would make the regression computationally expensive.

Our data consist of m triples of real numbers, in the form (P_k, Q_k, f^k) , the x-coordinate, the y-coordinate, and the intensity for location k , respectively. P , Q , and f are represented as row vectors. A non-standard reparametrization greatly simplified the calculations; this notation is borrowed from existing code which made use of those reparametrizations.

Each of the n Gaussians has

- an amplitude coefficient, a_i ,
- a variance coefficient, b_i ; actually $b_i = (1/2) \cdot (1/\text{variance})$,
- an abscissa (the x coordinate of the center of the Gaussian, x_i , and
- an ordinate (the y coordinate of the center of the Gaussian, y_i .

In `ggrad.m` and `ghess.m` (our gradient and Hessian computation code), a , b , x , and y are column vectors storing this information. These are our model parameters. For convenience, we define the vector

$$v = (a_1, \dots, a_n, b_1, \dots, b_n, x_1, \dots, x_n, y_1, \dots, y_n) \in \mathbb{R}^{4m} \quad (15)$$

containing all of our model parameters. We can now write our objective function, h , as

$$h(v) = \sum_{k=1}^m \left(\left(\sum_{i=1}^n g_i^k(v) \right) - f^k \right)^2 \quad (16)$$

where, g_i^k is the i th two-dimensional Gaussian evaluated at (P_k, Q_k) :

$$g_i^k = a_i \exp \left\{ -b_i \left((x_i - P_k)^2 + (y_i - Q_k)^2 \right) \right\}. \quad (17)$$

We seek the model parameter vector, v , which minimizes $h(v)$. Nonlinear regression is in general quite difficult. Here, we employ an iterative gradient-descent scheme to try to determine the optimal v , by starting with an initial guess and iterating by stepping along the negative gradient of h . The files `ggrad.m` and `ghess.m` contain MATLAB code for computing the gradient and Hessian, respectively, for any parameter vector.

4.4.2 Discussion

Location of Data Points. The data points should be all around the boundary of the masked regions, so we can identify all the individual Gaussian-like bumps. The data points must be chosen to be close to the masked object boundary since the scatter decreases very quickly away from the objects, and then noise can completely overwhelm the scatter. One way to do this is to start at a point on the boundary and move slightly in the outward normal direction, choose a good number of points in that area (not just one) – and then repeat around the entire boundary.

Noise. Noisy data are always problematic for regression problems, as one influential point can significantly affect the outcome. One way to address this problem is to apply a smoothing method in advance. It is important that only slight smoothing be done – so as to eliminate the most extreme outliers without significantly altering the good data. However, fitting data to a smooth curve is itself a form of smoothing. If we use sufficiently many data points, outliers will be less likely to have a noticeable impact. The obvious advantage of this approach is that there is no loss of information, since we are not altering the data by smoothing. The main disadvantage is that we need to use many more data points, and this may take more time than if we had simply smoothed the data in advance and used just one data point near each boundary point; also, it is still possible (though unlikely) that an outlier could still cause a significant problem (whether we smooth or not). Both of these approaches can work well, and it is also possible to employ a combination of them.

Model Adequacy and Future Work. The Compton scattering signature outside the masked region has tails which are thicker than those of Gaussians, indicating a slower decay rate. While using a sum of Gaussians makes good sense as a first attempt, since the overall shape of these functions is similar to the shape of the scatter, it would certainly be worthwhile to study this phenomenon to find a better model – a sum of thick-tailed Gaussian-like functions rather than a sum of Gaussians. Most of the work and code would be the same. Deriving the gradient and Hessian matrices for each new function set would be very involved, though we can make use of very accurate finite differencing methods.

4.4.3 Implementation

The implementation is not complete as of this writing.

4.5 Biharmonic Inpainting

Pete Schultz, Kevin Vixie, David Caraballo

4.5.1 Concept

We propose the following method to compute the contribution from scattering. The scattering field is known, except for random noise, in the regions of the image plane with no objects above them.

A reconstruction of the scattering fields in this region can be performed by isolating these regions, as in the previous method, and regularizing the data. In this way we can determine both the Dirichlet and Neumann boundary values at the edges of the objects' radiographic shadow. This suggests a fourth-order elliptic boundary-value problem. The biharmonic equation $\Delta^2 u = 0$ is a natural equation to try.

4.6 Tikhonov-like Regularization

Pete Schultz, Kevin Vixie

4.6.1 Concept

We propose the following method to compute the contribution from scattering. The noise-free scattering contribution is a smooth image, with positive values at every point. The Fourier coefficients corresponding to high frequencies should therefore be small; and therefore we expect

$$\int_{|\xi| \geq R} |\hat{u}(\xi)|^2$$

to be small if R is large.

On the other hand, the scattering field is known up to noise levels in the region corresponding to uninterrupted x-rays. Let d be the given radiographic data, and let B be the region of the image with empty space above it. Then $|u - d|$ should be small on B . By analogy with Tikhonov regularization, we seek to minimize

$$J(u) = \int_{|\xi| \geq R} |\hat{u}(\xi)|^2 + \lambda \int_B |u - d|^2.$$

5 Forward Measurement Operators

The goal of this subproject was to develop and demonstrate very fast methods for computing simulated radiographs of known objects of cylindrical or spherical symmetry. The time constraints are severe in that a full simulated radiograph must be computed in about one second. The emphasis is then on how to do the best possible approximations very rapidly. The emphasis is on incorporating nonlinear effects, though fast linear operators are also important.

5.1 Linear Polychromatic Approximation

Tom Asaki

5.1.1 Principles

I developed a simple linear approximation of a polychromatic forward measurement (projection) operator. This operator is nothing more than a multiple projection over a discrete energy spectrum. It ignores the effects of all scattered photons at the detector plane and photon energy changes due photon-matter interactions. That is, all photons that interact with material between source and detector are removed from consideration. The total energy R deposited at a radiograph pixel x is given by

$$R(x) = \int_E D(x, E) S(\Gamma, E) \exp \left\{ - \int_{\Gamma} \mu(s, E) ds \right\} dE. \quad (18)$$

The particle attenuation is given by the exponential of a line integral along path Γ from the source to detector location x of the material s and energy E dependent characteristic attenuation coefficient μ . S is a intensity source characterization multiplier that identifies the relative photon contributions along path Γ of energy E . D is a detector response function. We next consider the discrete 1-dimensional problem for objects of circular symmetry. The detector is naturally discrete for digitized radiographs. The object is divided into concentric rings, each with a constant material property description. Energy is also divided into several energy ranges. This problem can be formulated as a series of discrete Abel transforms:

$$R(x) = \sum_k D(x, E_k) S(\Gamma, E_k) \Delta E_k \exp \left\{ - P(s) \cdot \mu(E_k) \right\}. \quad (19)$$

The Abel matrix P is the transformation that takes a 2-dimensional radially symmetric object and projects its area onto 1-dimension. In effect, it is the discrete realization of the line integral operator of Eq. 18 for a parallel beam source.

5.1.2 Application

An example of energy transmission data is shown in Figure 23. The object is a set of nested spherical shells of material and outer radii: aluminum at 1.00 cm, air at 2.00 cm, iron at 3.00 cm. The figure reports distances from the object center in pixel units (1 pixel = 200 microns). The source is a discrete version of the Cygnus source S which has energy range $0 < E < 2.3$ MeV. The source is modeled in 0.1 MeV increments. The figure shows the transmission profile for each E_k as shades of blue; low energy responses in light blue and high energy responses in dark blue. The thick black line shows the sum R assuming a flat detector response. This plot illustrates the idea that no single value of μ can be an accurate description of a material.




Figure 23: Application of the linear polychromatic forward projection operator on an object of three layers imaged by the Cygnus source spectrum and a perfect energy-integrating detector. The blue lines show specific energy responses for $0 > E > 2.3$ MeV in increments of 0.1 MeV. Lighter lines correspond to lower energies. The thick black line is the total energy response.

This forward projection operator has been implemented in the cost functional evaluations used by the MVO reconstruction algorithms. It is used successfully in object reconstructions from simulated, but computationally and algorithmically independent, data. See section 6.1

5.2 Nonlinear Compton and Polychromatic Approximation

Rick Chartrand, Chris Orum, Patrick Campbell, David Dreisigmeyer

5.2.1 General problem statement

Our problem belongs in the general category of developing and demonstrating fast methods of computing simulated x-ray radiographs of known objects. More specifically, the problem is to develop a model of a forward measurement operator, an algorithm based on this model and its MATLAB implementation, that takes a given material description of an object, and returns a simulated radiograph, taking into account a polychromatic x-ray source and the non-linear effects of Compton scattering. The objects are assumed to be axially symmetric (invariant under rotation around a fixed axis) with special consideration for spherically symmetric objects (invariant under all rotations). A stated run-time objective is the ability to produce one simulated radiograph (for each material description) per second.

5.2.2 Working results

Our results are an algorithm, and MATLAB programs based on the algorithm. The algorithm for the MATLAB code `forward.m` is described in this subsection of this document. The development framework consists of:

1. a physics inspired model,
2. an algorithm based on the model,
3. MATLAB code implementing the algorithm.

Roughly, fidelity to physics propagates from 1 to 3 resulting in slower code; while aiming for faster code propagates from 3 to 1 resulting in a model less faithful to physics.

The output of the program `forward.m` coarsely approximates the x-ray radiograph of a spherically symmetric object. It returns a vector that is supposed to approximate the detector response along the 1-dimensional diameter of a rotationally invariant 2-dimensional detector response array. The run time of approximately six seconds is within an order of magnitude of the stated run time goal. It should be emphasized that this is for a large scenario of 50 energies and 6 materials. Simpler scenarios run faster. Equally important, this is for MATLAB code that is not fully optimized and running on a 1.3 MHz machine. The possibility of turning this into faster C code is a future project.

The ancillary program `makestuff.m` creates some of the inputs for `forward.m`.

5.2.3 Model descriptions

We describe two models for the forward measurement operator, called model A and model B. These are physics inspired but not physics exact. Model A is more realistic than model B, but the MATLAB implementation of an algorithm based on model A yielded code with an unacceptably slow running time, and therefore model B was taken as the working model.

Both models are amenable to implementation in either 2 or 3 dimensions, involving objects having either spherical or axial symmetry. The code `forward.m` is based on a model B in 2-dimensions involving spherically symmetric objects. We begin with a description of model A:

1. Assume an x-ray source consists of parallel rays whose intensity is given by $I_0(x, E)$, depending on position x and energy E . The source is polychromatic, meaning that E varies over a range of values and $I_0(x, E)$ generally varies with energy. While E belongs to a continuum in reality, we model E as belonging to a discrete set of at most about 50 different values.
2. The observed detector response $\text{radio}(x)$ is the sum over the product of a detector response function $R_D(x, E)$ and the intensity of the x-rays reaching the detector $I_D(x, E)$:

$$\text{radio}(x) = \sum_E R_D(x, E) I_D(x, E). \quad (20)$$

3. Photons that do not scatter contribute a linear term to $I_D(x, E)$, the x-ray intensity at the detector. The number of photons attenuate between the source and the detector:

$$\text{linear contribution to } I_D(x, E) = I_0(x, E) \exp \left\{ - \int_{\Gamma} \mu(s) ds \right\}. \quad (21)$$

Here Γ is the vertical line segment from source to detector, at position x , parametrized by s ; and $\mu(s)$ is a material-dependent attenuation coefficient. The attenuation coefficient incorporates both scattering and absorption.

4. The x-rays that do scatter, scatter exactly once between the source and the detector. The scattering angle distribution does not depend on the material, but the probability of scattering depends on the material. This is a macroscopic characterization of the material, captured by the parameter λ , the mean free path. For a layer of thickness Δs the probability that a photon, having arrived at that layer, scatters therein, is given by

$$\int_0^{\Delta s} \lambda^{-1} e^{-s/\lambda} ds = 1 - e^{-\Delta s/\lambda}. \quad (22)$$

5. Scattered photons are described by the energy dependent KN-probability distribution $Q(E, \theta)$ on the sphere. (KN = Klein-Nishina; this is *not* standard terminology). This is obtained from the Klein-Nishina formula, conditioned on the event that the photon scatters. The Klein-Nishina formula gives the probability of a single photon scattering off a single electron into a given solid angle. Since the integral of the function given by this formula over all angles is much less than one, this is normalized by dividing by the probability that the photon scatters. The notation for the 2-dimensional analog of $Q(E, \theta)$ is $q(E, \theta)$. (See equations (31), (34).)

6. Scattered x-rays attenuate between the source and the scattering point, as well as between the scattering point and the detector. Thus the non-linear component of $I_D(x, E)$ due to scattering may be obtained by integrating appropriately $I'_D(y \tan \theta, E)$ over θ and y where

$$I'_D(y \tan \theta, E) = \tag{23}$$

$$I_0(x, E) \exp \left\{ - \int_{\Gamma_1(y)} \mu(s) ds \right\} p_{\text{scat}}(y) q(E, \theta) \exp \left\{ - \int_{\Gamma_2(y, \theta)} \mu(s) ds \right\}.$$

Here $p_{\text{scat}}(y)$ is the probability of scattering at height y , $q(E, \theta)$ is the probability of scattering into angle θ (given that it scatters – the KN-probability), and the two exponential line integrals account for the pre- and post- scattering attenuation. See Figure 24.

7. The scattering process induces a change in energy in the x-rays that is related to the change in angle.




Figure 24: Source-object-detector geometry. For a scattering location at a given height y the angular scattering distribution as a function of θ may also be described as a function of detector position x . This is a pushforward measure under stereographic projection, where $\theta = \arctan(x/y)$. The line segment from x to the scattering point is denoted $\Gamma_1(y)$ and the line segment from the scattering point to the detector is denoted $\Gamma_2(y, \theta)$.

An attractive feature of one particular algorithmic implementation of this model (for objects with spherical symmetry, and in 2-dimensions) is that the two line integrals in (23) may computed by first parametrizing the scattering location via concentric circles, and then computing each as sums and differences of Abel transforms (a sum or difference depending on the position of the scattering point on the circle). Unfortunately, even with the inclusion

of this trick – which allows the use of previously defined code for computing the Abel transform – this leads to the unacceptably slow results mentioned above.

Pursuant to the stated run-time objective, model B was developed – and then adopted as the working model. Model B is obtained from model A by dropping items 6 and 7. Consequently, energy change is not incorporated into scattering, and the model has no post-scattering attenuation.

5.2.4 Algorithm description

A key reason for adopting model B over model A is that by ignoring post-scattering attenuation, the effect of Compton scattering may be computed by convolution.

For each energy E and each height y , let $S_{E,y}(x)$ denote the *scattering source intensity*. That is, having fixed E and y , $S_{E,y}(x)$ considered as a function of x , is the total intensity of the photons scattered at x (over all directions). This function $S_{E,y}(x)$ includes the source intensity at x , the effect of attenuation between the source and the scattering height y , and the probability of scattering at this position. Let $q_{E,y}(x)$ denote the angular scattering distribution as seen by the detector assuming the scattering occurs at the origin. This is a pushforward measure under stereographic projection; it is a sub-probability measure because the detector only sees a fraction of the total scattered photons. Then the contribution to the Compton scattering as seen by the detector (for energy E and height y) is the convolution

$$S_{E,y} * q_{E,y}(x). \quad (24)$$

The contribution to the non-linear portion of $I_D(x, E)$ due to Compton scattering is therefore

$$\text{non-linear contribution to } I_D(x, E) = \sum_E \sum_y S_{E,y} * q_{E,y}(x). \quad (25)$$

Equipped with this key idea, here are the main features of the program expressed in rough pseudo-code:

- For each energy E ,
 - Compute linear, unscattered flux (vectorized in x),
 - Compute $S_{E,y}(x)$ as given by equation (26), (vectorized in x, y),
 - Compute $S_{E,y} * q_{E,y}(x)$, (vectorized in x, y),
 - Sum over y to get the scattered flux (vectorized in x)
 - Compute $I_D(E, x)$ by adding unscattered and scattered flux,
 - Compute $R_D(E, x) \times I_D(E, x)$,
- Sum over E : compute $\text{radio}(x) = \sum_E R_D(x, E) I_D(x, E)$.

The scattering source intensity is computed as

$$S_{E,y}(x) = I_0(E, x)(1 - e^{-\Delta s/\lambda}) \exp \left\{ - \int_{\Gamma_1(y)} \mu(s) ds \right\}. \quad (26)$$

Here Δs arises from the discrete representation of the line segment $\Gamma_1(y)$ between the source and the scattering location. In our code, Δs is the same as dx , the pixel width.

5.2.5 The KN-scattering probability distribution

Here we describe the scattering probability distribution, its pushforward onto the detector, and how this measure on the detector may be represented for the purpose of computation.

We begin with the Wikipedia entry for the Klein-Nishina formula which provides a succinct and useful description of Compton scattering:

“The Klein-Nishina formula provides an accurate prediction of the angular distribution of x-rays and gamma-rays which are incident upon a single electron. The Klein-Nishina formula describes incoherent or Compton scatter.”

“More precisely, the Klein-Nishina formula provides the differential cross section with respect to solid angle of scatter, and it accounts for factors such as radiation pressure and relativistic quantum mechanics. For a photon of energy E , the differential cross section is:

$$\frac{d\sigma}{d\Omega} = 0.5r_e^2(P(E, \theta) - P(E, \theta)^2 \sin^2(\theta) + P(E, \theta)^3) \quad (27)$$

where θ is the angle of scatter; r_e is the classical electron radius; m_e is the mass of an electron; and $P(E, \theta)$ is the ratio of photon energy before and after the collision:

$$P(E, \theta) = \frac{1}{1 + (E/m_e c^2)(1 - \cos \theta)} \quad (28)$$

The value $d\sigma/d\Omega$ is the probability that a photon will scatter into the solid angle defined by $d\Omega = 2\pi \sin \theta d\theta$.” [4]

“In scattering, a differential cross section is defined by the probability to observe a scattered particle in a given quantum state per solid angle unit, such as within a given cone of observation, if the target is irradiated by a flux of one particle per surface unit.” [3]

$$\frac{d\sigma}{d\Omega} = \frac{\text{Scattered flux} / \text{Unit of solid angle}}{\text{Incident flux} / \text{Unit of surface}}. \quad (29)$$

The total probability that the photon scatters off the electron is the integral cross section

$$\sigma = \int_{4\pi} \frac{d\sigma}{d\Omega} d\Omega = \int_0^\pi p(E, \theta) 2\pi \sin \theta d\theta, \quad (30)$$

and we have introduced the function $p(E, \theta) = d\sigma/d\Omega$ as given by (27).

This integral is much less than 1, as it is for a single electron; it is material independent. The probability that the photon scatters at all in the material depends on the mean free path, which is a collective characterization of the material. Thus we make use of the conditional probability distribution

$$Q(E, \theta) \sin \theta d\theta d\phi = \frac{p(E, \theta) d\theta d\phi}{\int_0^{2\pi} \int_0^\pi p(E, \theta) \sin \theta d\theta d\phi}, \quad (31)$$

which we call the KN-probability distribution. In 3 dimensions, we model the scattered flux as the incident flux times the probability of scattering (depending on the mean free path λ) times this KN-probability that it scatters into a given solid angle.

For a 2-dimensional detector (3 dimensions) the *pushforward* measure in the plane of the detector describes the distribution of the conditional scattering probability as seen in the plane. It is a sub-probability (it integrates to something less than 1) because an infinite detector only sees scattering at angles less than $\pi/2$, and a finite detector sees even less. It is given by

$$Q_D(E, x_1, x_2)dx_1dx_2 = Q(E, \arctan(\frac{r}{y})) \frac{y}{(y^2 + r^2)^{3/2}} dx_1dx_2, \quad (32)$$

$$r = \sqrt{x_1^2 + x_2^2}$$

where the Jacobian $y/(y^2 + r^2)^{3/2}$ arises from the transformation of the surface element $\sin\theta d\theta d\phi$ through the projective transformation

$$\begin{cases} x_1 = y \sin\phi \tan\theta, \\ x_2 = y \cos\phi \tan\theta. \end{cases} \quad (33)$$

In 2-dimensions, the analog of (31) is the KN-probability

$$q(E, \theta)d\theta = \frac{\mathbf{p}(E, \theta)d\theta}{\int_{-\pi}^{\pi} \mathbf{p}(E, \theta)d\theta}, \quad (34)$$

and the pushforward measure on the line of the detector is

$$q_{E,y}(x)dx = q_D(E, x)dx = q(E, \arctan(\frac{x}{y})) \frac{y}{y^2 + x^2} dx. \quad (35)$$

Let $f(x)$ denote $q_{E,y}(x)$ with E and y fixed. The sub-probability density $f(x)$ on the detector is modeled as a discrete sum of point masses δ_x located over an array Λ of ‘pixels’ a width Δx apart. That is, the function $f(x)$ is approximated in a mathematical sense by

$$\sum_{x \in \Lambda} \left(\int_x^{x+\Delta x} f(x')dx' \right) \delta_x \approx \sum_{x \in \Lambda} f(x) \Delta x \delta_x. \quad (36)$$

Indeed, as $\Delta x \rightarrow 0$ the right hand side of (36) converges weakly to $f(x)$. Computationally however, we deal with this as vector of length

$$N = \lfloor \frac{\text{detector width}}{\Delta x} \rfloor \quad (37)$$

which is the column vector inside a multidimensional array

$$\text{scatangleprob}(y, :, E) = [f(\Delta x)\Delta x, f(2\Delta x)\Delta x, \dots, f(N\Delta x)\Delta x]', \quad (38)$$

where y and energy E are fixed.

5.2.6 Further Developments

The possibility of implementing this program in C is under consideration, as a means of making it faster. An analog of `forward.m` that deals with fully 3-dimensional objects (that is, not making a 2-dimensional approximation, but rather computing the Compton scattering as a 3-dimensional process creating a 2-dimensional simulated radiograph) would be slower, and the implementation in C of such an analog should also be investigated.

The magnitude of errors in `forward.m` driven by the goal of making fast running code needs to be quantified – in particular, this means considering the errors incurred by (a) dropping the post-scattering attenuation and (b) making the 2-dimensional approximation to the 3-dimensional process, and comparing the relative sizes of these errors.

5.2.7 Usage and Examples

To run `forward.m` it is first necessary to run `makestuff.m` which itself calls `MuMaker.m`. The function of `MuMaker.m` (written by Tom Asaki) is to create the attenuation coefficient matrix corresponding to categorical materials and discretized energies. Essentially whatever can be pre-processed is relegated to `makestuff.m`, leaving the essential computations in the forward measurement operator to `forward.m`.

Here is the usage of `forward.m`; underlined arguments come from `makestuff.m`:
`radio=forward(source, mu, lambda, radii, scatangleprob, dx, response, scatflag)`

Inputs:

- `source`: the 1-D intensity. Rows are energies, columns are source pixels, which are even in number.
- `mu`: the attenuation coefficients. Rows are energies, columns are materials.
- `lambda`: the mean free path for Compton scattering. It is assumed to be preprocessed so that rows are vertical layers of the object (assuming a vertical beam), columns are pixels that are covered by the object (which is assumed to be centered), and pages are energies.
- `radii`: the radii of the outer edge of each material, a row vector.
- `scatangleprob`: the pushforward of the scattering probability density to the detector (which is assumed to be the same size as the source). It is structured similarly to `lambda`, except the second dimension spans the whole source/detector width, and not just the object.
- `dx`: the length of the source/detector pixels.
- `response`: the detector acceptance function. Rows are energies, columns are detector pixels.
- `scatflag`: set to 0, this will cause scattering to be ignored.

Output: the 1-D radiograph.

As noted above, `makestuff.m` is a program that functions as a pre-processor. Here is its

usage:

`[mu,lambda,prob]=makestuff(radii,matvec,energies,dx,nmx,L)`

Object constructor: creates inputs for forward.m

Inputs:

- radii: vector of outer material radii. Must be same as vector given to forward.m.
- matvec: a vector of material numbers, corresponding to indices of the materials struct array created by MuMaker.
- energies: a vector of energies to be included in the spectrum. Should correspond to the rows of the vector source of forward.m
- dx: source pixel width. Same as given to forward.m.
- nmx: number of source pixels
- L: distance from detector to center of object

Outputs:

- mu: matrix of attenuation coefficients. Rows are energies, columns are radii.
- lambda: array of mean free paths for Compton scattering. Rows are object layers, columns are detector pixels (only those covered by object), pages are energies.
- prob: array giving scattering angle probability density, pushed forward to the detector. Pages are energies; each row is the detector density for a given depth in the object. These are only for determining the angle distribution; the likelihood of scatter is taken care of elsewhere.

Figures 25 through 27 show simulated radiographs of various objects composed of iron, aluminum, and air. The source is the polychromatic Cygnus source of uniform spatial intensity. The detector is 100% efficient and located 4 cm from the center of each object.

Figure 28 shows a final example that compares simulated radiographic signatures computed by two methods. The object is spherical and has four material layers (polyethylene, aluminum, air, and iron) with corresponding outer radii (1.749 cm, 2.427 cm, 3.082 cm, and 3.677 cm). The detector was assumed to have 100% efficiency at all energies and was placed 10 cm distant from the object center. The detector pixel size for the simulations was 200 microns. The source was a parallel beam of uniform intensity with the polychromatic Cygnus spectrum. Full particle transport calculations (MCNP) are shown as the noisy (rough and jagged) lines. The black line represents the attenuation only radiograph. The

red line represents the radiograph obtained by including the full MCNP physics package. Computations done using `forward.m` are shown as the clean (not rough or jagged) lines using the same color scheme. The attenuation only profiles are identical aside from noise. The profiles that include the scattering signature are most similar at large radii and differ the most towards the object center. This discrepancy might be due to the presence of Thompson scattering which is not modeled in the forward calculation. This calculation was performed without any adjustable parameters and takes into account a polychromatic source.




Figure 25: A 1-dimensional simulated radiograph of a 1 cm iron sphere.




Figure 26: A 1-dimensional simulated radiograph of a 1 cm inner-radius, 2 cm outer-radius aluminum spherical shell filled with air.




Figure 27: A 1-dimensional simulated radiograph of a 1 cm iron sphere nested within a 1 cm thick aluminum spherical shell.




Figure 28: Comparison between simulated radiographs computed by MCNP and by forward.m. See text for details.

5.3 Geometry and Harmonic Analysis

Chris Orum

5.3.1 Concept

Under the assumption that the x-rays scatter at most once between the source and the detector, for those x-rays that scatter we are faced with computing the attenuation integrals

$$\exp \left\{ - \int_{\Gamma_1(y)} \mu(s) ds \right\} \quad \text{and} \quad \exp \left\{ - \int_{\Gamma_2(y,\theta)} \mu(s) ds \right\}. \quad (39)$$

Figure 29 indicates how both of these may be computed with Abel transforms.

The Hankel, Fourier, and Abel transforms (H , F , and A respectively) are related by

$$H = FA,$$

or equivalently

$$A = F^{-1}H.$$

Our model contains a 2-dimensional object; so the Hankel transform is the zero-order Hankel transform, and the Fourier transform is 1-dimensional. (These transform relationships extend to higher dimensions.) The ability to express both pre- and post-scattering attenuation integrals with Abel transforms should be re-examined in light of these transform relationships. The idea is that as the scattering angle θ changes, the contribution to the photon intensity at the detector may be characterized as the previously mentioned Abel transforms shifted (depending on θ) and scaled (due to projection, and also depending on θ) and multiplied by Jacobian factors. If it turns out that these later three operations fit well with the Hankel and inverse Fourier transforms, then there is the possibility that the ensuing mathematical description of what the detector sees may be suitable to producing fast algorithms and code. The main point here is that a topic that needs more exploration is how harmonic analysis tools and geometry of our model are intertwined, and how they may be exploited.




Figure 29: After parameterizing the scattering locations with concentric circles, the pre- and post-scattering attenuation integrals may each be computed as the sum or difference of two Abel transforms that are determined by the outer and inner radii. Whether it is a sum or difference depends on where the scattering point occurs on the circle.

5.4 Fast Ray Casting

Matt Sottile

5.4.1 Concept

Another approach to approximating the radiograph of an object is to use algorithmic techniques from computer graphics that take advantage of algorithms and hardware developed within that field. Computer graphics programmers have required near real-time approximations for physical processes such as light passing through complex objects. There has been some work in algorithms inspired by such work in ray casting and ray tracking to build approximations for physical processes such as x-ray radiography. In addition, recent work on using specialized, ultra-high performance graphics processing units (GPUs) for implementing these algorithms within workstation class computers has proven to be promising. For Information, contact : Pat McCormick (CCS-1), John Turner (CCS-2).

5.5 Anomalous Diffusion

David Dreisigmeyer

5.5.1 Concept

When we consider the path of a single particle through an object the main feature is the low probability of a scattering event. This means that a particle will have long, uninterrupted flights between scattering events. Because of this, it seems likely that the processes under consideration will exhibit anomalous diffusion [5]. Specifically, we can reasonably expect super-diffusion to occur. Fortunately, this is an area attracting some attention recently by the mathematical community. Super-diffusion results in a fractional-ordered partial differential equation (FPDE) that interpolates between the heat and wave equations. The goal would be to find the Greens function for the resulting FPDE. This would allow for the rapid solution of the scattering problem. The main difficulty in this line of research is the complexity of the problem. We not only have long flights, but also changes in the energy of the particle, directionally dependent probabilities of scattering angles, and the possibility of particle absorption. All of these will affect the final form of the FPDE governing the problem.

6 Quantitative Object Reconstruction

The goal of this subproject was to demonstrate quantitative object reconstructions. A satisfactory reconstruction consists of shell layer thicknesses and material identifications. We wish to employ, as far as possible, new scattering-effects removal techniques and forward measurement operators, and examine how much improvement is obtained relative to our current methods. Our tool for quantitative reconstructions is a novel mixed-variable optimization method which we briefly describe.

6.1 Mixed-Variable Formulation

Tom Asaki, Mark Abramson, Kevin O'Reilly, John Dennis Jr.

6.1.1 Principles

The mixed-variable formulation of quantitative object reconstruction has been under development since the last quarter of FY05. Since the focus of this section is on reconstruction results, a full description of the method is not given here. However, because the method is relatively new and the application is novel, we present a brief overview and justification.

Mixed variable optimization (MVO) is a general scheme for functional optimization on parameter spaces that can include continuous, integer, and categorical variables along with constraints. Effective optimization must include new concepts of adjacency in parameter space (neighborhood) and efficient search strategies in the non-continuous variable subspace. The mathematics and implementation can be quite complex, but the potential (and now proven) benefits outweigh these difficulties. Of course we had the good fortune to contract the work with the premier applied mathematicians in the field (Mark Abramson and John E. Dennis Jr.), and build the subproject around a master's thesis (Kevin O'Reilly).

Consider the description of a non-dynamic cylindrically symmetric object of layered materials. Any slice perpendicular to the symmetry axis is circularly symmetric. Each slice is completely characterized by the number of material layers, the material within each layer, and the thickness of each layer. This description consists of an integer variable (number of layers), categorical variables (material within each layer), and continuous variables (thickness of each layer). The match with MVO is clear. The major benefit of the MVO framework is that we are able to implement, in a natural way, *all* of the prior knowledge we have about the object. This differs from regularized inversion techniques which favor reconstructions that have certain class properties (e.g. minimal total variation or smoothness characteristics). Instead, we can easily include very specific prior knowledge. Here is a list.

- Objects have a finite number of material boundaries. We may have lower and upper bounds on this number.
- Layers have a thickness. We may have upper or lower bounds on these thicknesses.
- Layers are composed of real materials. We can build a library of materials of interest and limit the reconstruction to those materials.
- Certain layers may be limited to a particular subset of our material library.
- Certain material pairs may never (or always) be adjacent.

There are additional MVO-related benefits. The computation provides the relative merits of additional solutions of differing number of layers and materials at no additional computation cost.

The MVO approach is not without its difficulties. The method will not necessarily converge to the global optimum of the complete parameter space. The computation typically

requires many hundreds of function evaluations, thus it is currently only suitable for use with very fast forward measurement operators.

6.1.2 Application

The current MVO code provides 1D object reconstructions from 1D flattened and normalized transmission data. At present the analysis process is user-intensive as we use a modified packaged MATLAB code and the steps to interface with this code are not yet automated. The current complexity of this interfacing is a direct result of the need to catalog and characterize the parameter space (including constraints, neighborhood description, and material library) and provide a cost function.

Typical MVO implementations require several hundred to a few thousand function evaluations for objects of two to five layers and a modest library of nine materials. Several factors affect how this translates into computational time, but about 1000 evaluations per minute is a good rule of thumb for my new laptop. The computational time scales roughly linearly with the number of layers (variable within the calculation) and quadratically with the library size (though careful neighborhood descriptions can reduce this to nearly linear).

Much work remains. Not all possible constraints are currently available. Results can be sensitive to initial guesses. The neighborhood structure needs to be carefully examined. Automated interfacing is needed.

6.2 Linear Polychromatic Reconstructions

Tom Asaki

6.2.1 Results

Now we present MVO reconstructions of ten standard test objects. Two radiographs were generated for each test case using the Los Alamos Monte Carlo transport code MCNP. The simulations assumed a parallel incident beam of uniform intensity with the Cygnus energy spectrum. The detector is an energy-integrating detector of 100% efficiency at all energies. The first set of radiographs is that obtained by the full physics package. The second set of radiographs excludes contributions from scattered photons – linear attenuation only. We demonstrate the use of a nontrivial forward measurement operator (see section 5.1) and examine the consequences of using a linear reconstruction process to model nonlinear effects.

We model the reconstruction process based on the assumptions of the second set of radiographs. We use the linear polychromatic forward measurement operator. We use a discrete approximation of the Cygnus source and coefficients interpolated from the NIST tables of X-ray attenuation properties of materials. We show reconstructions for all twenty data sets in Figures 30- 39. The lower graph in each figure shows the data with (red) and without (blue) the scattering signature. These simulations modeled 1 pixel as 200 microns. The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials are displayed within the subfigures as material abbreviations. Edge locations are displayed within the subfigures along the bargraph center as pixel values. The edge locations were rounded to the nearest integer pixel value.

Consider test object 3a (Figure 30). The simulated object is a set of nested spherical shells. The inner material is iron and has thickness 70 pixels (1.40 cm). This iron ball is surrounded by a Teflon shell of thickness 15 pixels (0.30 cm). This iron and Teflon sphere is surrounded by a third shell layer of iron of thickness 51 pixels (1.02 cm). This structure is seen as the left half of each of the two bar plots. The red curve in the data plot shows the simulated transmission profile of this object. the blue curve is the simulated transmission profile of the same object but with the photon signature of scattered photons removed. MVO applied to the blue data provides the object reconstruction labeled Rec-NS in the bar plot at left. The red data provides the object reconstruction labeled Rec-SC in the bar plot at right. In this case, the blue data is sufficient for an exact object reconstruction, both in terms of material identifications and layer thicknesses. The red data is insufficient for an exact reconstruction. Inner edge locations are misplaced and a beryllium layer replaces the Teflon layer. This is typical of reconstructions using data with uncorrected scattering or forward measurement operators that do not account for scattering. Reconstructions tend to favor lighter materials (less attenuating) and adjust edge locations to try and match a data profile that is poorly matched to the measurement operator.

The nine other test objects 3b-3j show similar reconstruction details. Some poor reconstructions do occur when using the data without scattering effects (blue). Figure 35 shows the common types of problems. Thin inner layers can be missed completely. Thin layers can be added. Edge locations can be soft. In addition, material identifications can be wrong, especially near the object center or for thin layers. This last problem occurs for object 3i (Figure 38) where an inner volume of air is determined to be polyethylene. All of these problems are due entirely to variations in the transmission signature caused by noise.

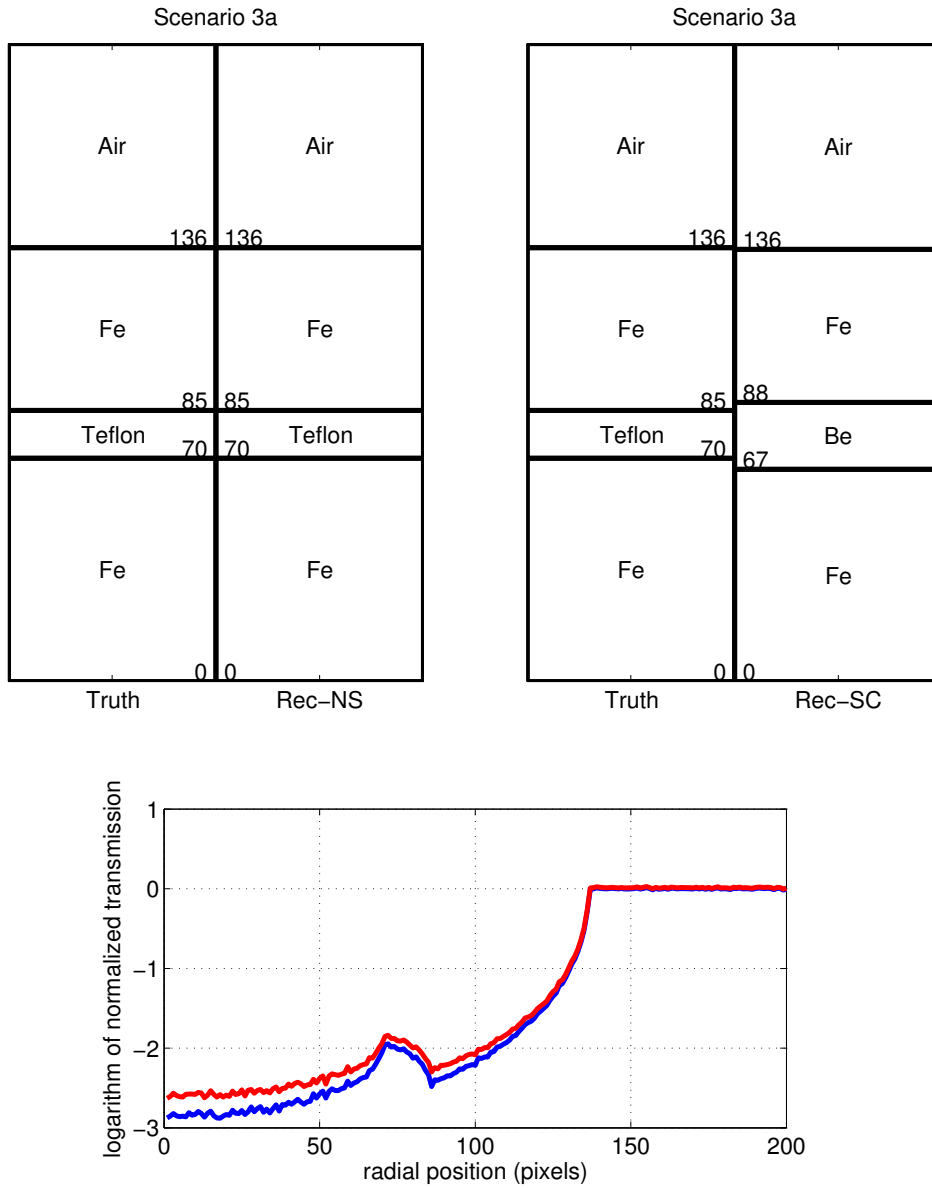


Figure 30: Simulated data and reconstructions of object 3a. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

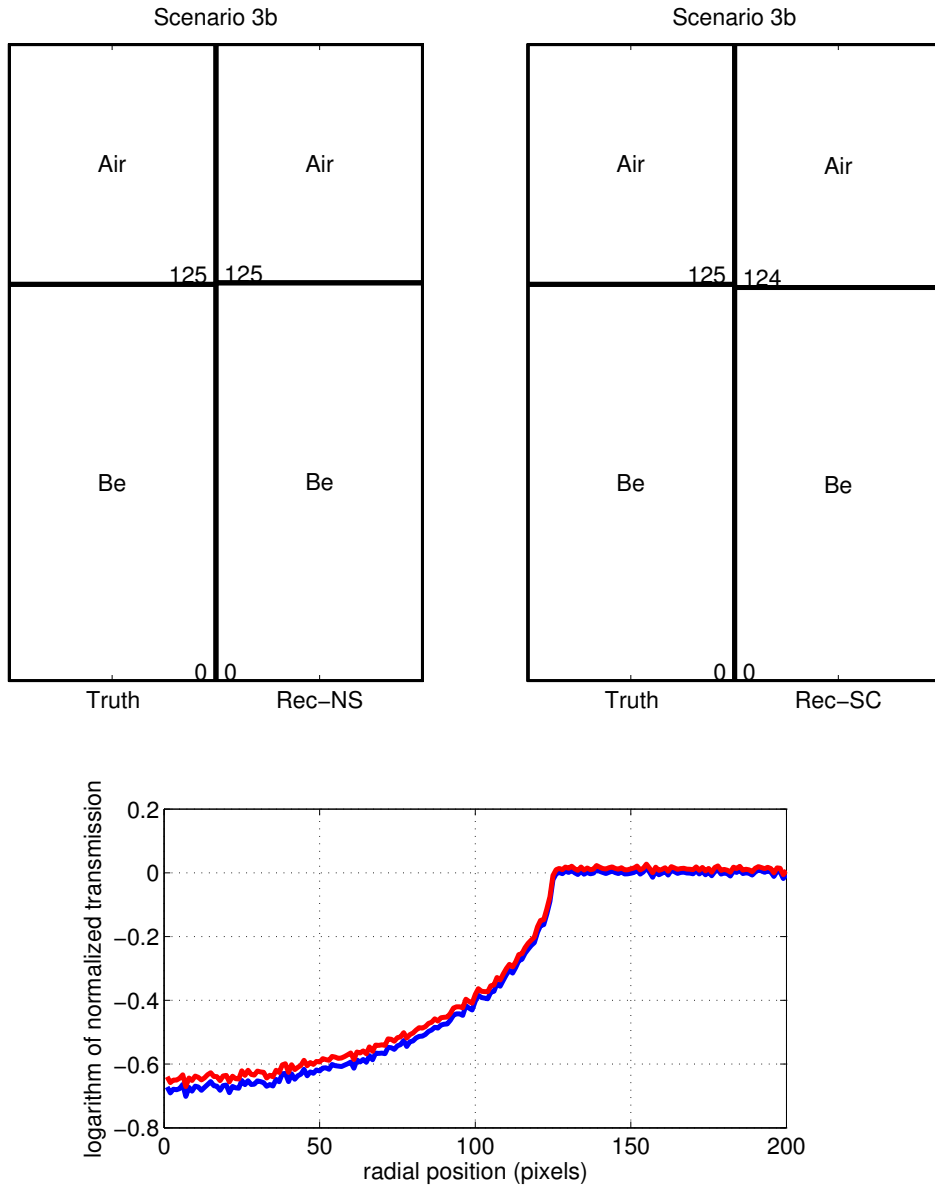


Figure 31: Simulated data and reconstructions of object 3b. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

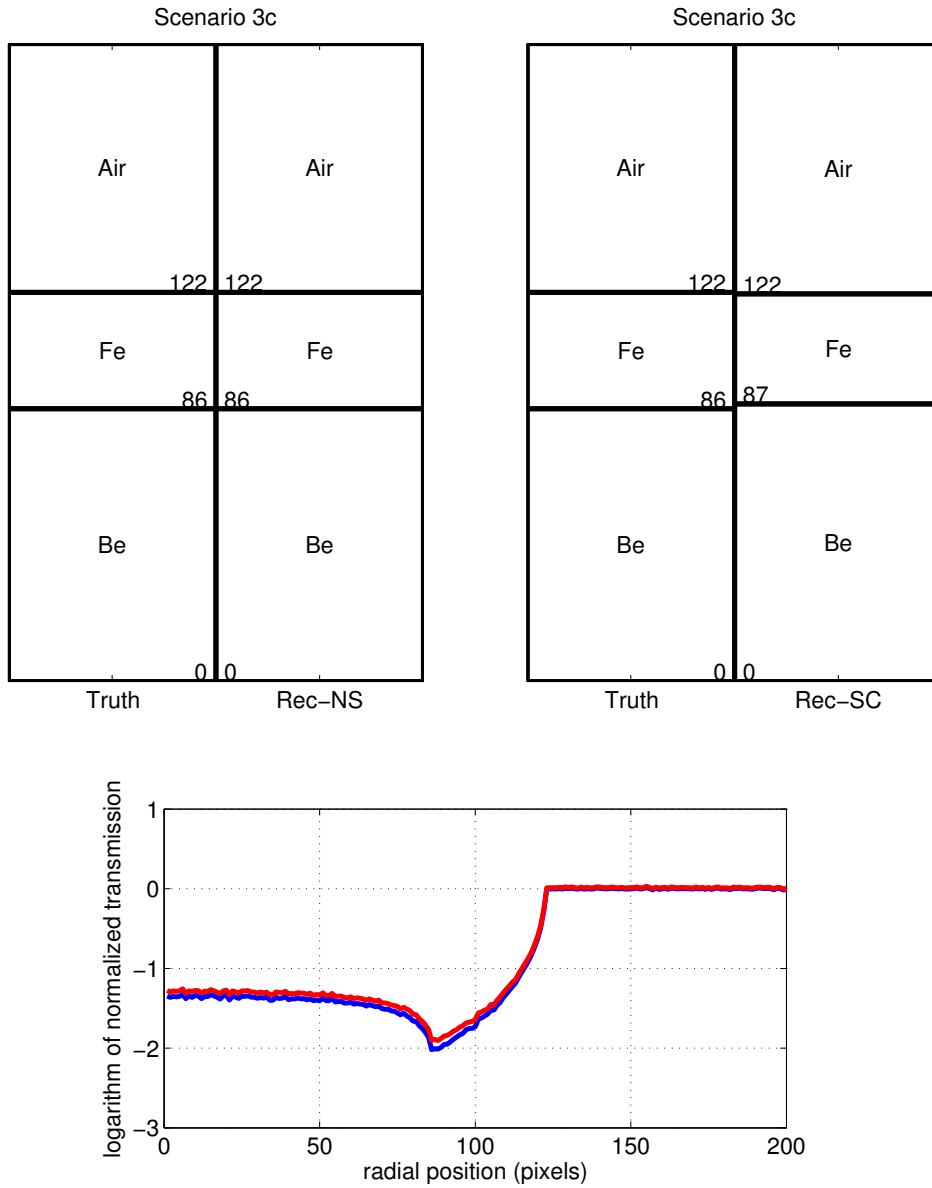


Figure 32: Simulated data and reconstructions of object 3c. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

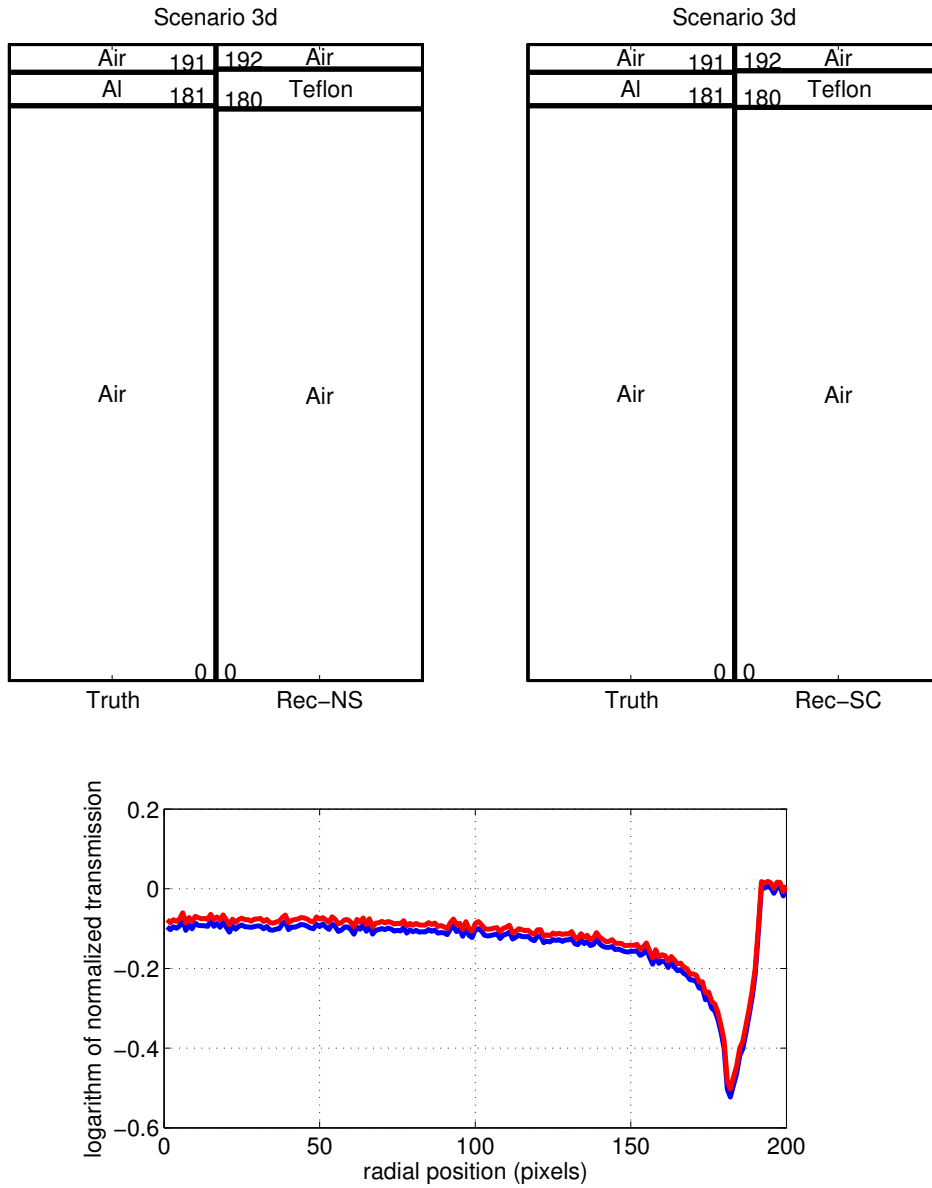


Figure 33: Simulated data and reconstructions of object 3d. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

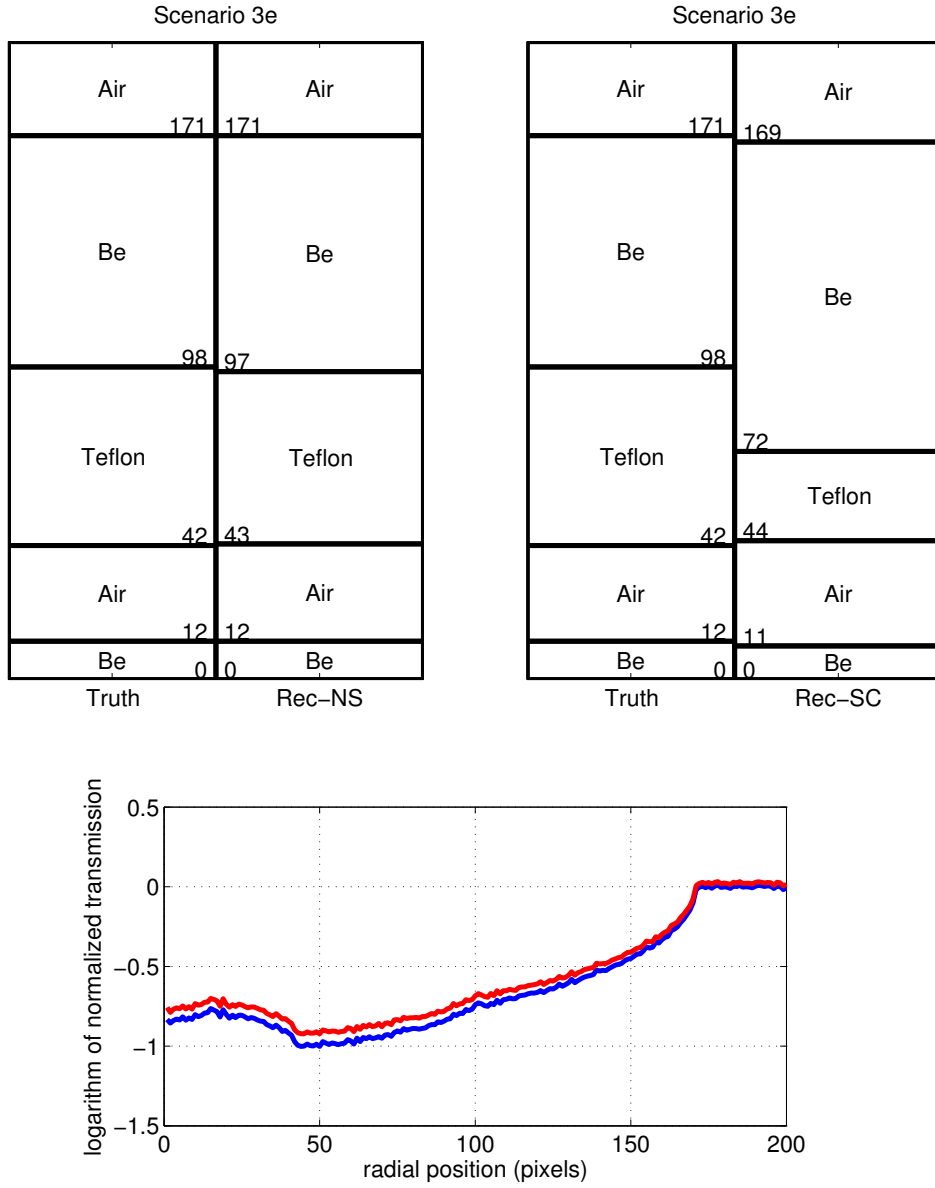


Figure 34: Simulated data and reconstructions of object 3e. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

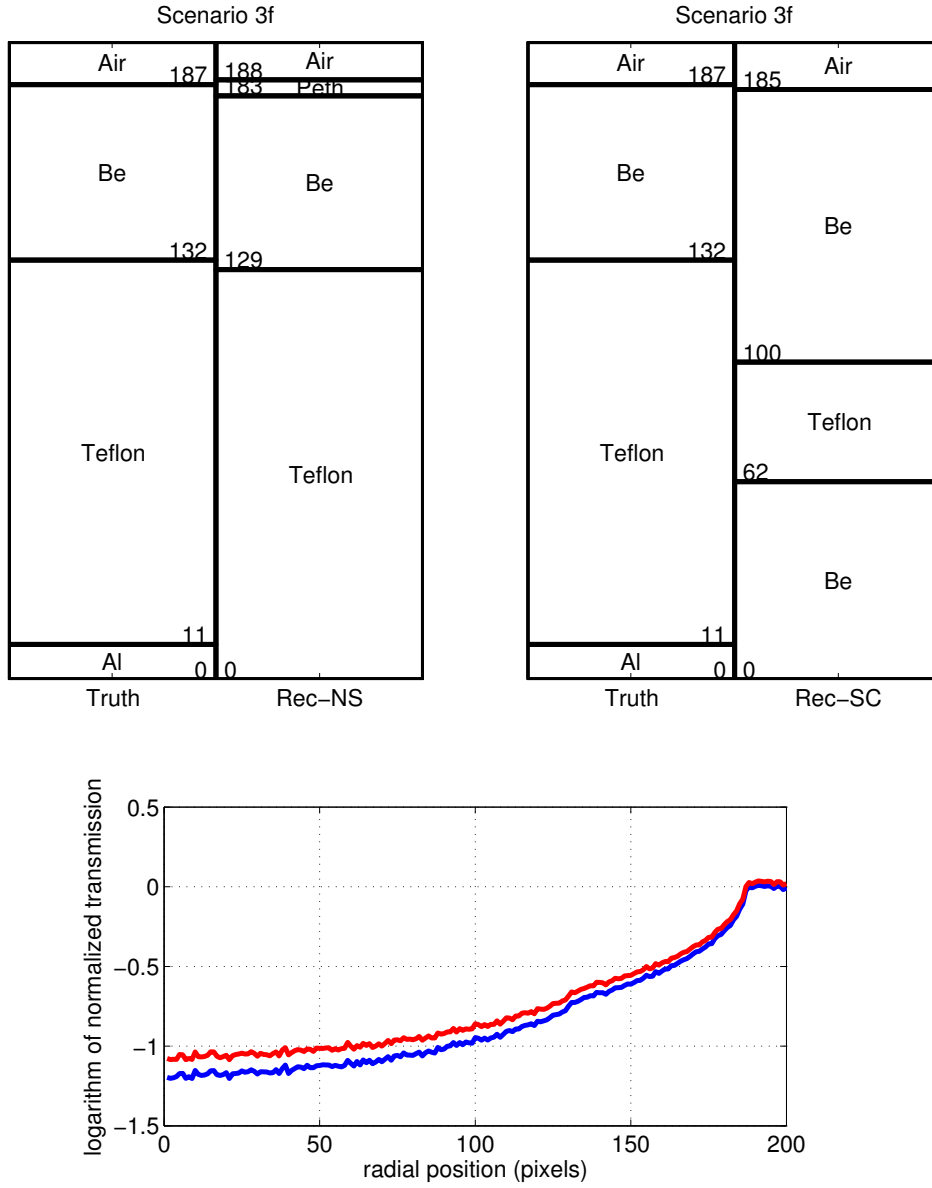


Figure 35: Simulated data and reconstructions of object 3f. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

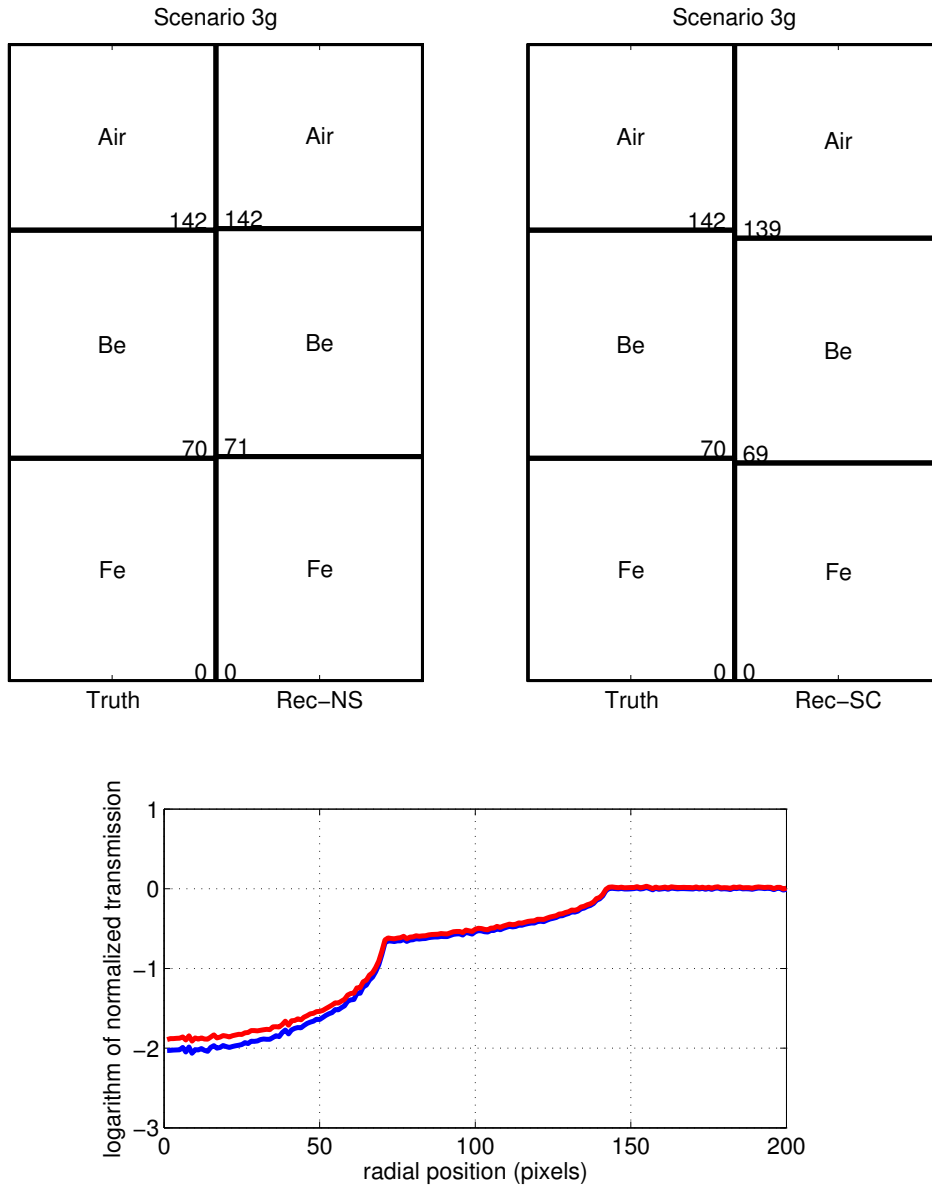


Figure 36: Simulated data and reconstructions of object 3g. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.


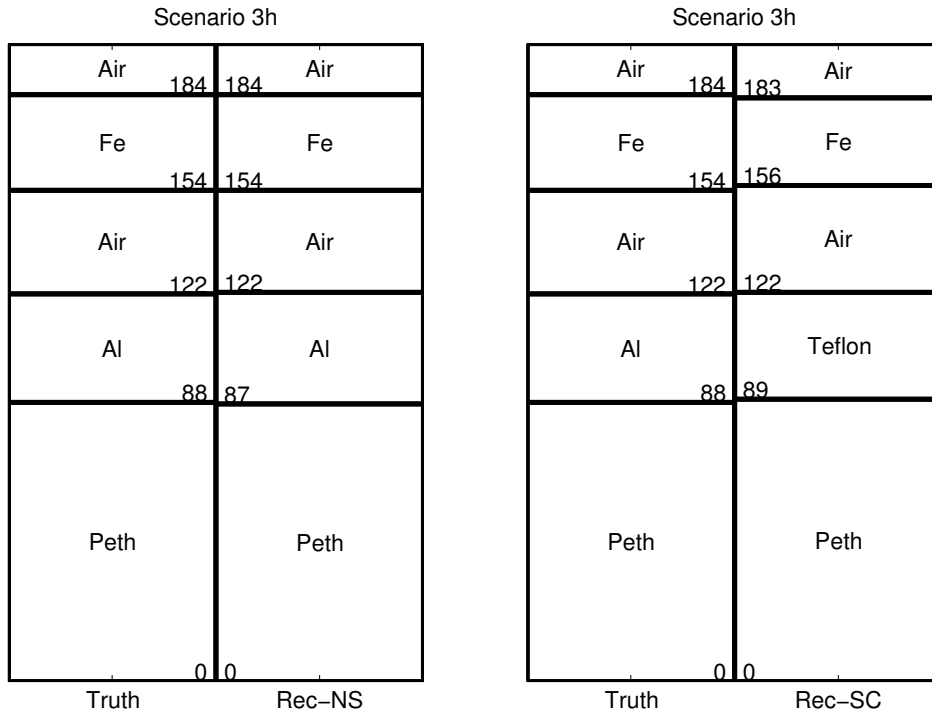


Figure 37: Simulated data and reconstructions of object 3h. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

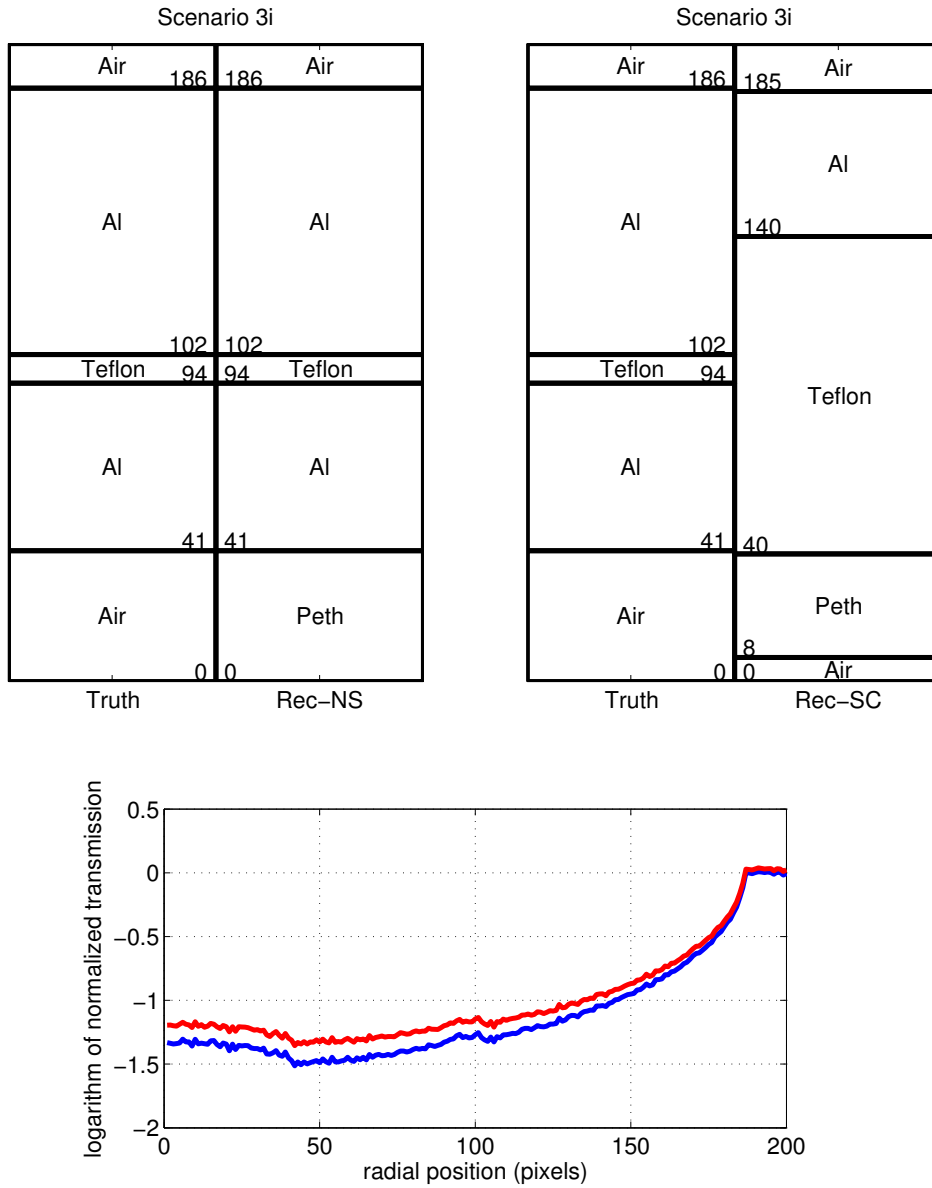


Figure 38: Simulated data and reconstructions of object 3i. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

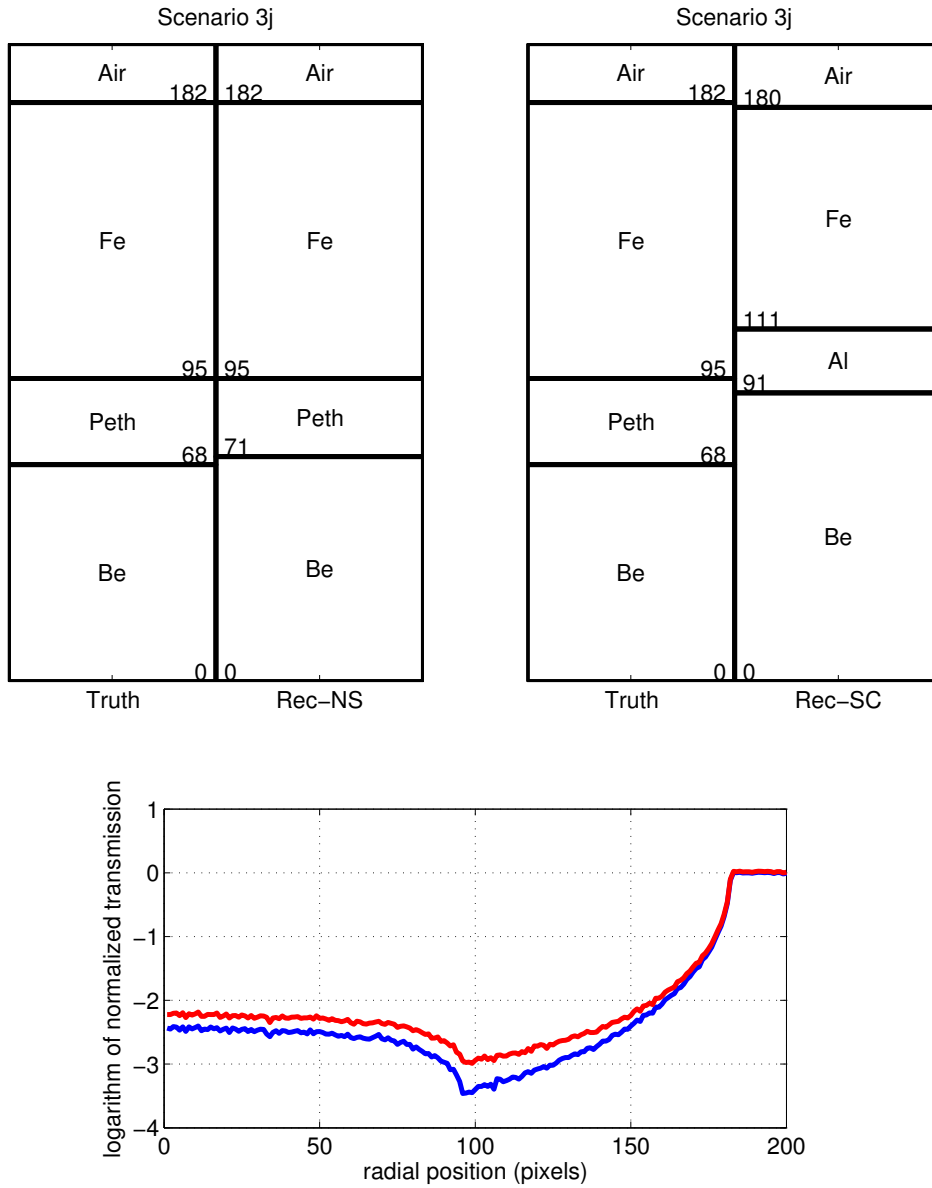


Figure 39: Simulated data and reconstructions of object 3j. The lower graph shows the data with (red) and without (blue) the scattering signature (1 pixel = 200 microns). The bar graphs show the true objects (left half of each subfigure) and reconstructions (right half of each subfigure) using data without (left subfigure) and with (right subfigure) scattering effects. Distance from the object center is shown from bottom to top. Materials and edge locations are displayed within the subfigures.

6.3 Reconstructions From Scattering-Effects Corrected Data

Tom Asaki

6.3.1 Results

The MVO method was used to compare reconstructions of object 3h from MCNP calculated data both with and without application of scattering-effects corrections. The results are shown in Figure 40. The simulation assumes a spatially uniform intensity in a parallel beam with the Cygnus source spectrum (Figure 23). The detector is located 10 cm from the object center and is assumed to be 100% efficient at all energies.

The data sets of interest are shown in Figure 40(f). The full MCNP computation provides the 1D radiograph shown as the red curve. If scattered photons are ignored, the result is the black curve. This is the curve most appropriate for reconstructions based on a linear measurement operator. The blue curve is a scattering effects corrected data set based on the 2-dimensional heat kernel regularization and inpainting method of section 4.2. The green curve is the scattering effects corrected data set based on the 1-dimensional Gaussian interpolation method of section 4.3.

The reconstructions are as follows.

- (a) This is a reconstruction from a noiseless radiograph that does not include contributions from scattered photons. This reconstruction is exact and represents the actual object, though edge locations are rounded to the nearest pixel for display purposes.
- (b) This reconstruction is from the black data. This is identical to case (a) except there is noise in the data. This example shows how much a Reconstruction can be affected by just the addition of noise. Material identifications are not affected but edge locations are less certain.
- (c) This reconstruction is from the red data. This example shows the effects of using raw data, uncorrected for the presence of a scattering signal, combined with a linear measurement operator. While edge locations are still fairly well determined in this case, there is a material misidentification. See section ?? for other examples.
- (d) This reconstruction is from the blue data. Here we see the first reconstruction on data modified in an attempt to remove the scattering signature. The result is quite good; it more closely approximates (a) than does (c).
- (e) This reconstruction is from the green data. This second example of corrected data and reconstruction also performs very well. It does the best job of capturing interior edge locations.

This example simulation was a fairly simple test. The scattering contribution is minimal and the variation in material thicknesses is small. Part of the reason for the stability in material identification is the use of a small material library (nine items). These results should be considered preliminary.

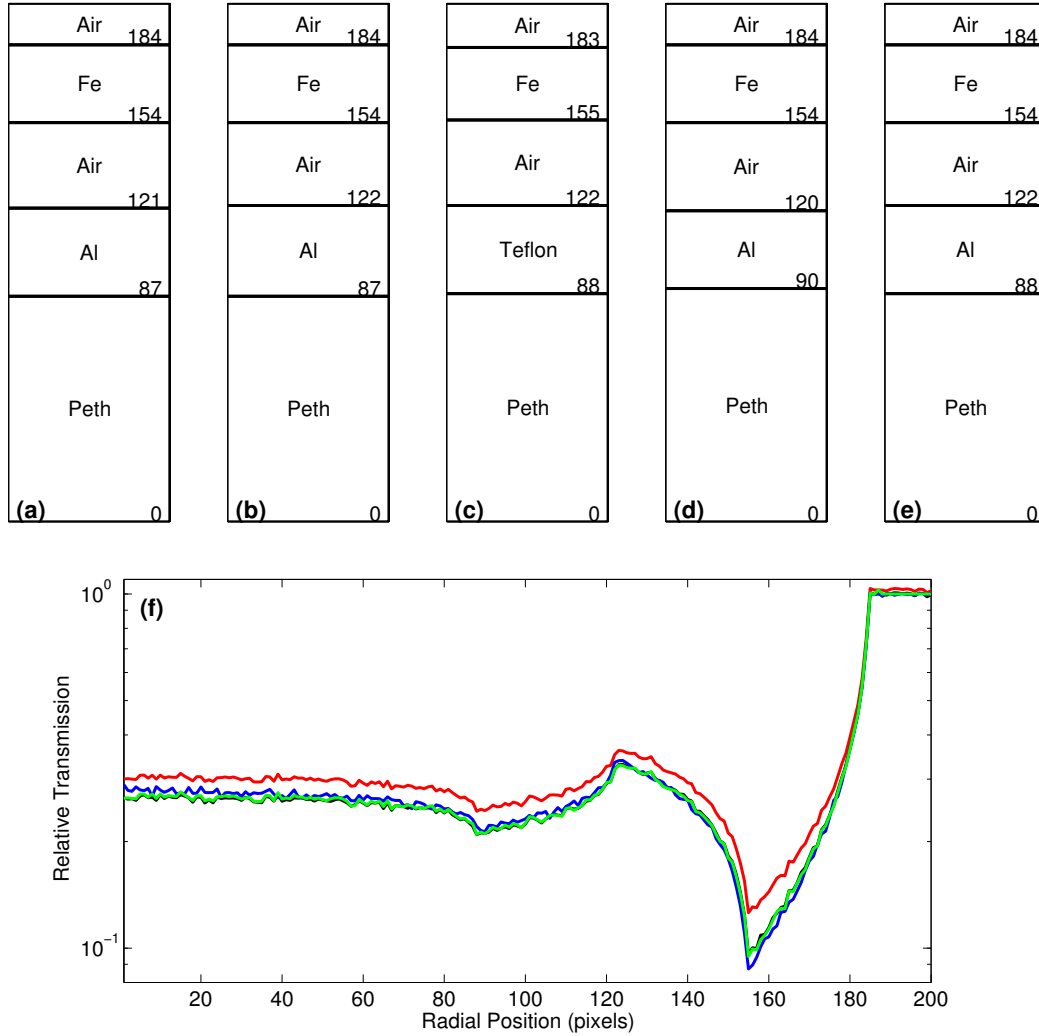


Figure 40: Simulated data and reconstructions of object 3h by several different methods. The lower graph (f) shows various computed and modified transmission radiograph (1 pixel = 200 microns). The bar graphs shows the corresponding reconstructions with distance from the object center shown from bottom to top. Materials and edge locations (rounded to nearest pixel) are displayed within the subfigures. The various reconstructions follow from the following data: (a) noiseless with no scattering (not shown data plot); (b) noisy data with no scattering (black); (c) noisy data with scattering (red); data corrected by HKI method (blue); data corrected by 1D-Gaussian method (green).

6.4 Nonlinear Polychromatic Reconstructions

Tom Asaki

6.4.1 Results

The nonlinear forward measurement operator (see section 5.2) was incorporated into a new cost function for MVO reconstructions based on raw data (no scattering effects corrections). One test has been completed. The chosen test object and simulation scenario was the same as that used in section 6.3. The result is shown in Figure 41. At left is the true object (identical to Figure 40(a)), and at right is the reconstruction. The edge locations are captured very well. However, the reduced scattering contribution corresponding to the object center (see Figure 28) leads to the same material misidentification as seen in the reconstruction using the linear measurement operator (see Figure 40(c)). It is anticipated that this problem is related to the use of relatively low source energies in these test scenarios. At higher energies, the Compton scattering assumptions of the forward model should be more accurate. These results should be considered preliminary.

Air	184
Fe	154
Air	121
Al	87
Peth	
(a)	0

Air	184
Fe	154
Air	122
Teflon	87
Peth	
(b)	0

Figure 41: Object and reconstruction example using the new fast forward measurement operator on raw radiographic data.

References

- [1] David W. Arathorn. *Map-Seeking Circuits in Visual Cognition: A Computational Mechanism for Biological and Machine Vision*. Stanford University Press, 2002.
- [2] T.F. Chan and Chiu-Kwong Wong. Total variation blind deconvolution. *IEEE Transactions on Image Processing*, 7:370 – 375, 1998.
- [3] http://en.wikipedia.org/wiki/Differential_cross_section.
- [4] http://en.wikipedia.org/wiki/Klein-Nishina_Formula.
- [5] Ralf Metzler and Joseph Klafter. The restaurant at the end of the random walk: recent developments in the description of anomalous transport by fractional dynamics. *Journal of Physics A: Mathematical and General*, 37:R161–R208, 2004.
- [6] Curtis R. Vogel. *Computational Methods for Inverse Problems*. SIAM, 2002.
- [7] Curtis R. Vogel, David Arathorn, Albert Parker, and Austin Roorda. Retinal motion tracking in adaptive optics scanning laser ophthalmoscopy. 2005.